

# FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM Running on VMware and Hyper-V

Deployment Guide for FlexPod Hosting SQL Server 2017 Databases Running on RHEL7.4 in VMware ESXi 6.7 and Windows Server 2016 Hyper-V Virtual Environments

Last Updated: January 25, 2019



# About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2019 Cisco Systems, Inc. All rights reserved.

# Table of Contents

Executive Summary .....	10
Solution Overview .....	11
Introduction.....	11
Audience .....	11
Purpose of this Document.....	11
Highlights of this Solution .....	11
Technology Overview .....	12
FlexPod System Overview .....	12
FlexPod Benefits .....	12
FlexPod: Cisco and NetApp Verified and Validated Architecture .....	12
Integrated System .....	13
Out-of-the-Box Infrastructure High Availability .....	13
FlexPod Design Principles .....	13
Cisco Unified Computing System.....	13
Cisco UCS Manager .....	14
Cisco UCS Fabric Interconnects .....	15
Cisco UCS 5108 Blade Server Chassis.....	15
Cisco UCS B200 M5 Blade Server .....	16
Cisco UCS Fabric Extenders.....	16
Cisco VIC Interface Cards.....	17
Cisco UCS Differentiators .....	17
Cisco Nexus 93180YC-EX Switches .....	19
NetApp AFF A300 Storage .....	19
NetApp ONTAP 9.5.....	20
NetApp SnapCenter .....	21
SnapCenter Architecture .....	21
SnapCenter Features .....	22
VMware vSphere 6.7 .....	24
Microsoft Windows Server 2016 Hyper-V.....	24
Red Hat Enterprise Linux 7.....	24
Microsoft SQL Server 2017 .....	24
Solution Design.....	25
Deployment Hardware and Software.....	27
Software Revisions.....	27
Configuration Guidelines.....	27
Physical Infrastructure .....	29

FlexPod Cabling .....	29
Network Switch Configuration .....	31
Physical Connectivity .....	31
FlexPod Cisco Nexus Base .....	31
Set Up Initial Configuration .....	31
Cisco Nexus A .....	31
Cisco Nexus B .....	32
FlexPod Cisco Nexus Switch Configuration .....	33
Enable Licenses .....	33
Set Global Configurations .....	33
Create VLANs .....	34
Add NTP Distribution Interface .....	35
Add Individual Port Descriptions for Troubleshooting .....	35
Create Port Channels .....	37
Configure Port Channel Parameters .....	38
Configure Virtual Port Channels .....	39
Uplink into Existing Network Infrastructure .....	40
Storage Configuration .....	41
NetApp Hardware Universe .....	41
Controllers and Disk Shelves .....	41
AFF A300 Controllers .....	41
Disk Shelves .....	41
Clustered Data ONTAP 9.5 .....	41
Complete Configuration Worksheet .....	41
Configure ONTAP Nodes .....	41
Infrastructure Storage Virtual Machine (Optional) .....	57
Create Storage Virtual Machine .....	57
Create Load-Sharing Mirrors of SVM Root Volume .....	57
Create Block Protocol (iSCSI) Service .....	58
Configure HTTPS Access .....	58
Configure NFSv3 .....	59
Create FlexVol Volumes .....	59
Create Boot LUNs .....	60
Schedule Deduplication .....	60
Create Management LIF .....	60
Create iSCSI LIFs .....	60
Create NFS LIF .....	60

Add Infrastructure SVM Administrator .....	61
Gather Necessary Information .....	61
SVM for SQL Linux Virtual Machines on ESXi .....	62
Create Storage Virtual Machines .....	62
Create Load-Sharing Mirrors of SVM Root Volume .....	62
Create Block Protocol (iSCSI) Service .....	63
Configure HTTPS Access .....	63
Create FlexVol Volumes .....	64
Create ESXi Boot LUNs .....	64
Create Storage Volume for SQL Virtual Machine Datastore and Swap LUNs .....	64
Create SQL Data and Log LUNs .....	65
Schedule Deduplication .....	65
Create iSCSI LIFs .....	65
Add SVM Administrator .....	66
Gather Necessary Information .....	66
SVM for SQL Linux Virtual Machines on Hyper-V .....	67
Create Storage Virtual Machine .....	67
Create Load-Sharing Mirrors of SVM Root Volume .....	67
Create Block Protocol (iSCSI) Service .....	68
Configure HTTPS Access .....	68
Create FlexVol Volumes .....	69
Create Hyper-V Boot LUNs .....	69
Create Storage Volume for SQL Virtual Machines Running on Hyper-V .....	69
Create SQL Data and Log LUNs .....	70
Schedule Deduplication .....	70
Create iSCSI LIFs .....	70
Add SVM Administrator .....	70
Gather Necessary Information .....	71
Server Configuration .....	73
Cisco UCS Base Configuration .....	73
Perform Initial Setup of Cisco UCS Fabric Interconnects for FlexPod Environments .....	73
Cisco UCS Setup .....	74
Log in to Cisco UCS Manager .....	74
Upgrade Cisco UCS Manager Software to Version 4.0(1c) .....	74
Anonymous Reporting .....	74
Configure Cisco UCS Call Home .....	75
Add Block of IP Addresses for KVM Access .....	75

Synchronize Cisco UCS to NTP .....	76
Edit Chassis Discovery Policy .....	77
Enable Server and Uplink Ports .....	77
Acknowledge Cisco UCS Chassis and FEX .....	78
Create Uplink Port Channels to Cisco Nexus Switches .....	78
Create Organization (Optional) .....	80
Create MAC Address Pools .....	80
Create IQN Pools for iSCSI Boot .....	82
Create IP Pools for iSCSI Boot .....	83
Create UUID Suffix Pool .....	84
Create Server Pool .....	84
Create VLANs .....	85
Modify Default Host Firmware Package .....	88
Set Jumbo Frames in Cisco UCS Fabric .....	89
Create Local Disk Configuration Policy (Optional) .....	89
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) .....	90
Create Power Control Policy .....	91
Create Server BIOS Policy .....	92
Create Maintenance Policy .....	96
Create vNIC Templates .....	97
Creating Adapter Policy .....	101
Create LAN Connectivity Policy for iSCSI Boot .....	102
Create vMedia Policy for installing both ESXi and Hyper-V Hypervisors .....	105
Create Boot Policy for iSCSI Boot .....	107
Create Service Profile Templates .....	109
Create vMedia-Enabled Service Profile Template .....	116
Create Service Profiles .....	117
Add More Servers to FlexPod Unit .....	117
Gather Necessary Information .....	118
ONTAP Boot, Data LUNs, and Igroups Setup .....	119
Create igroups .....	119
VMware Cluster .....	119
Map Boot LUNs to igroups .....	119
VMware vSphere 6.7U1 Setup .....	121
VMware ESXi 6.7U1 .....	121
Download Cisco Custom Image for ESXi 6.7U1 .....	121
Launching KVM Console Server .....	121

Set Up VMware ESXi Installation .....	121
Install ESXi.....	122
Set Up Management Networking for ESXi Hosts.....	122
Verifying UEFI secure boot of ESXi.....	124
Log into VMware ESXi Hosts by Using VMware Host Client.....	125
Set Up Basic VMkernel Ports on Standard vSwitch .....	125
Setup iSCSI Multipathing .....	127
Add Required Datastores .....	128
Configure NTP on ESXi Hosts .....	129
VMware vCenter 6.7 .....	130
Build the VMware vCenter Server Appliance .....	130
Set Up VMware vCenter Server .....	134
Configure the Default Swap File Location .....	137
Configure the ESXi Power Management .....	137
Add AD User Authentication to vCenter (Optional) .....	137
ESXi Dump Collector Setup for iSCSI-Booted Hosts .....	139
Cisco UCS Manager Plug-in for VMware vSphere Web Client.....	140
FlexPod VMware vSphere Distributed Switch (vDS) .....	140
Configure the VMware vDS in vCenter .....	140
Creating Virtual Machine for SQL Server.....	151
Microsoft Windows Server 2016 Hyper-V Deployment .....	154
Install Microsoft Windows Server 2016.....	154
Install Windows Server 2016 .....	155
Install Chipset Driver Software .....	156
Install Windows Roles and Features.....	157
Configure Networking .....	158
Install NetApp Host Utilities .....	160
Configure Multipath-IO.....	161
Host Renaming and Join to Domain .....	162
Configure Microsoft iSCSI Initiator.....	163
Create Windows Failover Cluster .....	168
Prepare the Shared Disks for Cluster Use .....	168
Test and Validate the Cluster .....	170
Create a New Failover Cluster .....	171
Post Cluster Creation Tasks.....	172
Create Clustered Virtual Machine for RHEL .....	175
RHEL on Virtual Machine Installation and Configuration .....	180

Download Red Hat Enterprise Linux 7.4 image .....	180
Install Red Hat Enterprise Linux .....	180
RHEL OS Configuration for Higher Database Performance .....	181
Join Linux Virtual Machine to Microsoft Windows Active Directory (AD) Domain .....	182
Configure iSCSI Network Adapters for Storage Access .....	183
Install and Configure NetApp Linux Unified Host Utilities, Software iSCSI Initiator and Multipath .....	185
SQL Server 2017 on RHEL Installation and Configuration .....	196
SQL Server 2017 Installation .....	196
Change Memory Settings .....	197
Change the Default Directories for Databases .....	197
Configure Tempdb Database Files .....	198
Add SQL Server to Active Directory Domain .....	198
SQL Server Always On Availability Group on RHEL Virtual Machines Deployment .....	200
Install and Configure Pacemaker Cluster on RHEL Virtual Machines .....	200
Install and Configure Always On Availability Group on Pacemaker Cluster .....	204
Automatic Failover Testing of Availability Group .....	212
FlexPod Management Tools Setup .....	214
Deploy NetApp Virtual Storage Console 7.2.1 .....	214
Virtual Storage Console 7.2.1 Pre-installation Considerations .....	214
Install Virtual Storage Console 7.2.1 .....	214
Register Virtual Storage Console with vCenter Server .....	219
Enable VASA and SRA .....	220
Discover and Add Storage Resources .....	222
Optimal Storage Settings for ESXi Hosts .....	224
Virtual Storage Console 7.2.1 Provisioning Datastores .....	225
Provision NFS Datastore .....	225
Provision iSCSI Datastore .....	229
Deploy NetApp SnapCenter .....	231
SnapCenter Server Requirements .....	231
SnapCenter 4.1.1 License Requirements .....	233
SnapCenter Server .....	233
SnapCenter Plug-in for VMware vSphere .....	233
Support for Virtualized Databases and File Systems .....	234
SnapCenter Installation .....	234
SnapCenter Configuration .....	234
Install SnapCenter Plug-in for VMware .....	235
Host and Privilege Requirements for the Plug-in for VMware vSphere .....	235

Run As Credentials.....	236
Install the Plug-in for VMware vSphere from the SnapCenter GUI.....	237
Configure SnapCenter Plug-in for vCenter.....	237
View Virtual Machine Backups and Restore from vCenter by Using SnapCenter Plug-in.....	244
View Backups.....	244
Restore from vCenter by Using SnapCenter Plug-in.....	246
Summary .....	250
Appendix.....	251
FlexPod Backups.....	251
Cisco UCS Backup.....	251
Cisco Nexus Backups .....	252
About the Authors.....	254
Acknowledgements.....	254



## Executive Summary

---

It is important that a datacenter solution embrace technology advancement in various areas, such as compute, network, and storage technologies to address rapidly changing requirements and challenges of IT organizations.

FlexPod is a popular converged datacenter solution created by Cisco and NetApp. It is designed to incorporate and support a wide variety of technologies and products into the solution portfolio. There have been continuous efforts to incorporate advancements in the technologies into the FlexPod solution. This enables the FlexPod solution to offer a more robust, flexible, and resilient platform to host a wide range of enterprise applications.

This document discusses a FlexPod reference architecture using the latest hardware and software products and provides deployment recommendations for hosting Microsoft SQL Server databases in VMware ESXi and Microsoft Windows Hyper-V virtualized environments, with Linux support enablement from Microsoft for SQL Server deployment.

The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release 4.0.1c to support the Cisco UCS hardware platforms including Cisco UCS B-Series Blade Servers, Cisco UCS 6300 Fabric Interconnects, Cisco Nexus 9000 Series Switches, and NetApp AFF Series Storage Arrays.

# Solution Overview

---

## Introduction

The current IT industry is witnessing vast transformations in the datacenter solutions. In the recent years, there is a considerable interest towards pre-validated and engineered datacenter solutions. Introduction of virtualization technology in the key areas has impacted the design principles and architectures of these solutions in a big way. It has opened the doors for many applications running on bare metal systems to migrate to these new virtualized integrated solutions.

FlexPod System is one such pre-validated and engineered datacenter solution designed to address rapidly changing needs of IT organizations. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed compute, network and storage components to serve as the foundation for a variety of enterprise workloads including databases, ERP, CRM and Web applications, etc.

With Microsoft SQL Server 2017, Microsoft has made a big announcement to support SQL Server deployments on Linux Operating systems. SQL Server 2017 on Linux platforms brings in support for most of the major features that are currently supported by SQL in a Windows platform. Microsoft claims that database performance of SQL on Linux should be similar as that of SQL in Windows. SQL on Linux is as secured as SQL in Windows platform. High availability features such as Always On Availability Groups and Failover Cluster Instance are well supported and integrated in Linux operating systems.

This enablement relieves Windows centric deployments of SQL Server and offers flexibility to customers to choose and deploy SQL Server databases on variety of widely used Linux operating systems.

## Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

## Purpose of this Document

This document discusses reference architecture and step-by-step deployment guidelines for deploying Microsoft SQL Server 2017 databases on Red Hat Enterprise Linux Operating system on FlexPod system. It also provides deployment steps for configuring SQL Always On Availability Group feature for achieving high availability of SQL databases.

## Highlights of this Solution

The following software and hardware products distinguish the reference architecture from others.

- SQL Server 2017 deployment on Red Hat Enterprise Linux 7.4
- SQL Always On Availability Group configuration for high availability of databases
- Cisco UCS B200 M5 Blade Servers
- NetApp All Flash A300 storage with Data ONTAP 9.5 and NetApp SnapCenter 4.1.1 for virtual machine backup and recovery
- 40G end-to-end networking and storage connectivity
- VMWare vSphere 6.7 and Windows Server 2016 Hyper-V

## Technology Overview

---

### FlexPod System Overview

FlexPod is a best practice datacenter architecture that includes these components:

- Cisco Unified Computing System
- Cisco Nexus Switches
- NetApp FAS or AFF storage, NetApp E-Series storage systems

These components are connected and configured according to best practices of both Cisco and NetApp and provide the ideal platform for running multiple enterprise workloads with confidence. The reference architecture covered in this document leverages the Cisco Nexus 9000 Series switch. One of the key benefits of FlexPod is the ability to maintain consistency at scaling, including scale up and scale out. Each of the component families shown in Figure 7 (Cisco Unified Computing System, Cisco Nexus, and NetApp storage systems) offers platform and resource options to scale the infrastructure up or down, while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

### FlexPod Benefits

As customers transition toward shared infrastructure or cloud computing they face several challenges such as initial transition hiccups, return on investment (ROI) analysis, infrastructure management and future growth plan. The FlexPod architecture is designed to help with proven guidance and measurable value. By introducing standardization, FlexPod helps customers mitigate the risk and uncertainty involved in planning, designing, and implementing a new datacenter infrastructure. The result is a more predictive and adaptable architecture capable of meeting and exceeding customers' IT demands.

The following list provides the unique features and benefits that FlexPod system provides for consolidation SQL Server database deployments.

- Support for latest Intel Xeon processor scalable family CPUs, UCS B200 M5 blades enables consolidating more SQL Server virtual machines and thereby achieving higher consolidation ratios reducing Total Cost of Ownership and achieving quick ROIs.
- End to End 40 Gbps networking connectivity using Cisco third-generation fabric interconnects, Nexus 9000 series switches and NetApp AFF A300 storage Arrays.
- Blazing IO performance using NetApp All Flash Storage Arrays and Complete virtual machine protection by using NetApp Snapshot technology and direct storage access to SQL virtual machines using in-guest iSCSI Initiator.
- Nondisruptive policy-based management of infrastructure using Cisco UCS Manager.

### FlexPod: Cisco and NetApp Verified and Validated Architecture

Cisco and NetApp have thoroughly validated and verified the FlexPod solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their datacenters to this shared infrastructure model. This portfolio includes, but is not limited to the following items:

- Best practice architectural design

- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for FlexPod configuration do's and don'ts)
- Frequently asked questions (FAQs)
- Cisco Validated Designs and NetApp Verified Architectures (NVAs) focused on many use cases

Cisco and NetApp have also built a robust and experienced support team focused on FlexPod solutions, from customer account and technical sales representatives to professional services and technical support engineers. The cooperative support program extended by Cisco and NetApp provides customers and channel service partners with direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues. FlexPod supports tight integration with virtualized and cloud infrastructures, making it a logical choice for long-term investment. The following IT initiatives are addressed by the FlexPod solution.

## Integrated System

FlexPod is a pre-validated infrastructure that brings together compute, storage, and network to simplify, accelerate, and minimize the risk associated with datacenter builds and application rollouts. These integrated systems provide a standardized approach in the datacenter that facilitates staff expertise, application onboarding, and automation as well as operational efficiencies relating to compliance and certification.

## Out-of-the-Box Infrastructure High Availability

FlexPod is a highly available and scalable infrastructure that IT can evolve over time to support multiple physical and virtual application workloads. FlexPod has no single point of failure at any level, from the server through the network, to the storage. The fabric is fully redundant and scalable, and provides seamless traffic failover, should any individual component fail at the physical or virtual layer.

## FlexPod Design Principles

FlexPod addresses four primary design principles:

- Application availability: Makes sure that services are accessible and ready to use
- Scalability: Addresses increasing demands with appropriate resources
- Flexibility: Provides new services or recovers resources without requiring infrastructure modifications
- Manageability: Facilitates efficient infrastructure operations through open standards and APIs

The following sections provide a brief introduction of the various hardware and software components used in this solution.

## Cisco Unified Computing System

Cisco Unified Computing System is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless 40 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Cisco Unified Computing System accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non-virtualized systems. Cisco Unified Computing System provides:

- Comprehensive Management
- Radical Simplification
- High Performance

Cisco Unified Computing System consists of the following components:

- Compute - The system is based on an entirely new class of computing system that incorporates rack mount and blade servers based on Intel Xeon scalable processors product family.
- Network - The system is integrated onto a low-latency, lossless, 40-Gbps unified network fabric. This network foundation consolidates Local Area Networks (LANs), Storage Area Networks (SANs), and high-performance computing networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- Virtualization - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access - The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. It is also an ideal system for Software defined Storage (SDS). Combining the benefits of single framework to manage both the compute and Storage servers in a single pane, Quality of Service (QOS) can be implemented if needed to inject IO throttling in the system. In addition, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity. In addition to external storage, both rack and blade servers have internal storage which can be accessed through built-in hardware RAID controllers. With storage profile and disk configuration policy configured in Cisco UCS Manager, storage needs for the host OS and application data gets fulfilled by user defined RAID groups for high availability and better performance.
- Management - the system uniquely integrates all system components to enable the entire solution to be managed as a single entity by Cisco UCS Manager. Cisco UCS Manager has an intuitive GUI, a CLI, and a powerful scripting library module for Microsoft PowerShell built on a robust API to manage all system configuration and operations.

Cisco Unified Computing System fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

## Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, embedded management for all software and hardware components in the Cisco UCS. Using Single Connect technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive GUI, CLI, or an XML API. Cisco UCS Manager resides on a pair of Cisco UCS 6300 Series Fabric Interconnects using a clustered, active-standby configuration for high availability.

Cisco UCS Manager offers unified embedded management interface that integrates server, network, and storage. Cisco UCS Manager performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers comprehensive set of XML API for third part integration, exposes 9000 points of integration and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Service profiles benefit both virtualized and non-virtualized environments and increase the mobility of non-virtualized servers, such as when moving workloads from server to server or taking a server offline for service or upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.

For more information about Cisco UCS Manager, go to: <https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

## Cisco UCS Fabric Interconnects

The Fabric interconnects provide a single point for connectivity and management for the entire system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency regardless of a server or virtual machine's topological location in the system.

Cisco UCS 6300 Series Fabric Interconnects support the bandwidth up to 2.43-Tbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic using a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blade, rack servers, and virtual machines are interconnected using the same mechanisms. The Cisco UCS 6332-16UP is a 1-RU Fabric Interconnect that features up to 40 universal ports that can support 24 40-Gigabit Ethernet, Fiber Channel over Ethernet, or native Fiber Channel connectivity. In addition to this it supports up to 16 1- and 10-Gbps FCoE or 4-, 8- and 16-Gbps Fibre Channel unified ports.

**Figure 1 Cisco UCS Fabric Interconnect 6332-16UP**



For more information, go to: <https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6332-16up-fabric-interconnect/index.html>

## Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis. The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors. Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2304 Fabric Extenders. A passive mid-plane provides multiple 40 Gigabit Ethernet connections between blade servers and fabric interconnects. The Cisco UCS 2304 Fabric Extender has four 40 Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 can provide one 40 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total eight 40G interfaces to the compute. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.


For more information, go to: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html>

## Cisco UCS B200 M5 Blade Server

The Cisco UCS B200 M5 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA. The Cisco UCS B200 M5 Blade Server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager Software and simplified server access through Cisco SingleConnect technology. The Cisco UCS B200 M5 Blade Server is a half-width blade. Up to eight servers can reside in the 6-Rack-Unit (6RU) Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit of blade chassis in the industry. You can configure the Cisco UCS B200 M5 to meet your local storage requirements without having to buy, power, and cool components that you do not need. The Cisco UCS B200 M5 Blade Server provides these main features:

- Up to two Intel Xeon Scalable CPUs with up to 28 cores per CPU
- 24 DIMM slots for industry-standard DDR4 memory at speeds up to 2666 MHz, with up to 3 TB of total memory when using 128-GB DIMMs
- Modular LAN On Motherboard (mLOM) card with Cisco UCS Virtual Interface Card (VIC) 1340, a 2-port, 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable mLOM mezzanine adapter
- Optional rear mezzanine VIC with two 40-Gbps unified I/O ports or two sets of 4 x 10-Gbps unified I/O ports, delivering 80 Gbps to the server; adapts to either 10- or 40-Gbps fabric connections
- Two optional, hot-pluggable, Hard-Disk Drives (HDDs), Solid-State Disks (SSDs), or NVMe 2.5-inch drives with a choice of enterprise-class RAID or pass-through controllers

**Figure 2 Cisco UCS B200 M5 Blade Server**



For more information, go to: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/datasheet-c78-739296.html>

## Cisco UCS Fabric Extenders

Cisco UCS 2304 Fabric Extender brings the unified fabric into the blade server enclosure, providing multiple 40 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management. It is a third-generation I/O Module (IOM) that shares the same form factor as the second-generation Cisco UCS 2200 Series Fabric Extenders and is backward compatible with the shipping Cisco UCS 5108 Blade Server Chassis. The Cisco UCS 2304 connects the I/O fabric between the Cisco UCS 6300 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fibre Channel over Ethernet (FCoE) fabric to connect all blades and chassis together. Because the fabric extender is similar to a distributed line card, it does not perform any switching and is managed as an extension of the fabric interconnects. This approach reduces the overall infrastructure complexity and enabling Cisco UCS to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

The Cisco UCS 2304 Fabric Extender has four 40Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 can

provide one 40 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total eight 40G interfaces to the compute. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.

**Figure 3 Cisco UCS 2304 Fabric Extender**



For more information, go to: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-675243.html>

## Cisco VIC Interface Cards

The Cisco UCS Virtual Interface Card (VIC) 1340 is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fiber Channel over Ethernet (FCoE) capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. All the blade servers for both Controllers and Computes will have MLOM VIC 1340 card. Each blade will have a capacity of 40Gb of network traffic. The underlying network interfaces will share this MLOM card.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs).

For more information, go to: <http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.html>

## Cisco UCS Differentiators

Cisco Unified Computing System is revolutionizing the way servers are managed in the data center. The following are the unique differentiators of Cisco UCS and Cisco UCS Manager:

- **Embedded Management**—In Cisco UCS, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
- **Unified Fabric**—In Cisco UCS, from blade server chassis or rack servers to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters which in turn reduce capital and operational expenses of the overall solution.

- Auto Discovery—By simply inserting the blade server in the chassis or connecting rack server to the fabric interconnect, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where its compute capability can be extended easily while keeping the existing external connectivity to LAN, SAN and management networks.
- Policy Based Resource Classification—When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy-based resource classification of Cisco UCS Manager.
- Combined Rack and Blade Server Management—Cisco UCS Manager can manage B-Series blade servers and C-Series rack server under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
- Model-based Management Architecture—Cisco UCS Manager Architecture and management database is model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, Templates—The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Loose Referential Integrity—In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each-other. This provides great flexibility where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.
- Policy Resolution—In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
- Service Profiles and Stateless Computing—a service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in Multi-Tenancy Support—The combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and a service profiles based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- Extended Memory—the enterprise-class Cisco UCS B200 M5 blade server extends the capabilities of Cisco’s Unified Computing System portfolio in a half-width blade form factor. The Cisco UCS B200 M5 harnesses the power of the latest Intel Xeon scalable processors product family CPUs with up to 3 TB of RAM— allowing huge virtual machine-to-physical server ratio required in many deployments or allowing large memory operations required by certain architectures like Big-Data.
- Virtualization Aware Network—VM-FEX technology makes the access network layer aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when


virtual network is managed by port-profiles defined by the network administrators' team. VM-FEX also off-loads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.

- Simplified QoS—Even though Fiber Channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

## Cisco Nexus 93180YC-EX Switches

In this solution, the Cisco Nexus 93180YC-EX Switches are used as upstream switches. It offers 48 1/10/25/40/100 Gigabit Ethernet SFP+ ports and 6 40/100-Gbps QSFP+ uplink ports. All ports are line rate, delivering 3.6 Tbps of throughput with a latency of less than 2 micro seconds in a 1-rack-unit (1RU) form factor.

**Figure 4** Cisco UCS Nexus 93180YC-EX




For more information, go to: <https://www.cisco.com/c/en/us/products/switches/nexus-93180yc-ex-switch/index.html>

## NetApp AFF A300 Storage

With the new A-Series All Flash FAS (AFF) controller lineup, NetApp provides industry leading performance while continuing to provide a full suite of enterprise-grade data management and data protection features. The A-Series lineup offers double the IOPS, while decreasing the latency.

This solution utilizes the NetApp AFF A300. This controller provides the high performance benefits of 40GbE and all flash SSDs, while taking up only 5U of rack space. Configured with 24 x 3.8TB SSD, the A300 provides ample performance and over 60TB effective capacity. This makes it an ideal controller for a shared workload converged infrastructure. For situations where more performance is needed, the A700s would be an ideal fit.

**Figure 5** NetApp AFF A300



## NetApp ONTAP 9.5

ONTAP 9.5 is the data management software that is used with the NetApp all-flash storage platforms in the solution design. ONTAP software offers unified storage for applications that read and write data over block or file-access protocols in storage configurations that range from high-speed flash to lower-priced spinning media or cloud-based object storage.

ONTAP implementations can run on NetApp engineered FAS or AFF appliances, on commodity hardware (ONTAP Select), and in private, public, or hybrid clouds (NetApp Private Storage and Cloud Volumes ONTAP). Specialized implementations offer best-in-class converged infrastructure as featured here as part of the FlexPod Datacenter solution and access to third-party storage arrays (NetApp FlexArray virtualization).

Together these implementations form the basic framework of the NetApp Data Fabric, with a common software-defined approach to data management and fast, efficient replication across platforms. FlexPod and ONTAP architectures can serve as the foundation for both hybrid cloud and private cloud designs.

The following few sections provide an overview of how ONTAP 9.5 is an industry-leading data management software architected on the principles of software defined storage.

### NetApp Storage Virtual Machine

A NetApp ONTAP cluster serves data through at least one and possibly multiple storage virtual machines (SVMs; formerly called Vservers). An SVM is a logical abstraction that represents the set of physical resources of the cluster. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and might reside on any node in the cluster to which the SVM has been given access. An SVM might own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node in the storage cluster to another. For example, a NetApp FlexVol flexible volume can be non-disruptively moved to a new node and aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and thus it is not tied to any specific physical hardware.

An SVM can support multiple data protocols concurrently. Volumes within the SVM can be joined together to form a single NAS namespace, which makes all of an SVM's data available through a single share or mount point to NFS and CIFS clients. SVMs also support block-based protocols, and LUNs can be created and exported by using iSCSI, FC, or FCoE. Any or all of these data protocols can be configured for use within a given SVM. Storage administrators and management roles can also be associated with SVM, which enables higher security and access control, particularly in environments with more than one SVM, when the storage is configured to provide services to different groups or set of workloads.

### Storage Efficiencies

Storage efficiency has always been a primary architectural design point of ONTAP data management software. A wide array of features allows you to store more data using less space. In addition to deduplication and compression, you can store your data more efficiently by using features such as unified storage, multitenancy, thin provisioning, and utilize NetApp Snapshot technology.

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduce the total logical capacity used to store customer data by 75 percent, a data reduction ratio of 4:1. This space reduction is enabled by a combination of several different technologies, including deduplication, compression, and compaction.

Compaction, which was introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the ONTAP WAFL file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk to save space.

## Encryption

Data security continues to be an important consideration for customers purchasing storage systems. NetApp supported self-encrypting drives in storage clusters prior to ONTAP 9. However, in ONTAP 9, the encryption capabilities of ONTAP are extended by adding an Onboard Key Manager (OKM). The OKM generates and stores keys for each of the drives in ONTAP, allowing ONTAP to provide all functionality required for encryption out of the box. Through this functionality, sensitive data stored on disk is secure and can only be accessed by ONTAP.

Beginning with ONTAP 9.1, NetApp has extended the encryption capabilities further with NetApp Volume Encryption (NVE), a software-based mechanism for encrypting data. It allows a user to encrypt data at the per-volume level instead of requiring encryption of all data in the cluster, thereby providing more flexibility and granularity to the ONTAP administrators. This encryption extends to Snapshot copies and NetApp FlexClone volumes that are created in the cluster. One benefit of NVE is that it runs after the implementation of the storage efficiency features, and, therefore, it does not interfere with the ability of ONTAP to create space savings.

For more information about encryption in ONTAP, see the [NetApp Power Encryption Guide](#) in the [NetApp ONTAP 9 Documentation Center](#).

## FlexClone

NetApp FlexClone technology enables instantaneous cloning of a dataset without consuming any additional storage until cloned data differs from the original.

## SnapMirror (Data Replication)

NetApp SnapMirror is a replication technology for data replication across different sites, or within the same datacenter, or on-the-premises datacenter to cloud, or cloud to on-the-premises datacenter.

## NetApp SnapCenter

SnapCenter is a NetApp next-generation data protection software for tier 1 enterprise applications. SnapCenter, with its single-pane-of-glass management interface, automates and simplifies the manual, complex, and time-consuming processes associated with the backup, recovery, and cloning of multiple databases and other application workloads.

SnapCenter leverages technologies, including NetApp Snapshot copies, SnapMirror replication technology, SnapRestore data recovery software, and FlexClone thin cloning technology, that allow it to integrate seamlessly with technologies offered by Oracle, Microsoft, SAP, VMware, and MongoDB across FC, iSCSI, and NAS protocols. This integration allows IT organizations to scale their storage infrastructure, meet increasingly stringent SLA commitments, and improve the productivity of administrators across the enterprise.


SnapCenter is used in this solution for backup and restore of VMware virtual machines. SnapCenter has the capability to backup and restore SQL databases running on Windows OS. SnapCenter currently does not support SQL Server running on Linux OS, but that functionality is coming in near future. SnapCenter does not support the backup and restore of virtual machines running on Hyper-V. For Hyper-V virtual machine backup, NetApp [SnapManager for Hyper-V](#) can be used.

## SnapCenter Architecture

SnapCenter is a centrally managed web-based application that runs on a Windows platform and remotely manages multiple servers that must be protected.

Figure 6 illustrates the high-level architecture of the NetApp SnapCenter Server.

Figure 6 SnapCenter Architecture



The SnapCenter Server has an HTML5-based GUI as well as PowerShell cmdlets and APIs.

The SnapCenter Server is high-availability capable out of the box, meaning that if one SnapCenter host is ever unavailable for any reason, then the second SnapCenter Server can seamlessly take over and no operations are affected.

The SnapCenter Server can push out plug-ins to remote hosts. These plug-ins are used to interact with an application, a database, or a file system. In most cases, the plug-ins must be present on the remote host so that application- or database-level commands can be issued from the same host where the application or database is running.

To manage the plug-ins and the interaction between the SnapCenter Server and the plug-in host, SnapCenter uses SM Service, which is a NetApp SnapManager web service running on top of Windows Server Internet Information Services (IIS) on the SnapCenter Server. SM Service takes all client requests such as backup, restore, clone, and so on.

The SnapCenter Server communicates those requests to SMCORE, which is a service that runs co-located within the SnapCenter Server and remote servers and plays a significant role in coordinating with the SnapCenter plug-ins package for Windows. The package includes the SnapCenter plug-in for Microsoft Windows Server and SnapCenter plug-in for Microsoft SQL Server to discover the host file system, gather database metadata, quiesce and thaw, and manage the SQL Server database during backup, restore, clone, and verification.

SnapCenter Virtualization (SCV) is another plug-in that manages virtual servers running on VMWare and that helps in discovering the host file system, databases on virtual machine disks (VMDK), and raw device mapping (RDM).

## SnapCenter Features

SnapCenter enables you to create application-consistent Snapshot copies and to complete data protection operations, including Snapshot copy-based backup, clone, restore, and backup verification operations. SnapCenter provides a centralized management environment, while using role-based access control (RBAC) to

delegate data protection and management capabilities to individual application users across the SnapCenter Server and Windows hosts.

SnapCenter includes the following key features:

- A unified and scalable platform across applications and database environments and virtual and nonvirtual storage, powered by the SnapCenter Server
- Consistency of features and procedures across plug-ins and environments, supported by the SnapCenter user interface
- RBAC for security and centralized role delegation
- Application-consistent Snapshot copy management, restore, clone, and backup verification support from both primary and secondary destinations (NetApp SnapMirror and SnapVault)
- Remote package installation from the SnapCenter GUI
- Nondisruptive, remote upgrades
- A dedicated SnapCenter repository for faster data retrieval
- Load balancing implemented by using Microsoft Windows network load balancing (NLB) and application request routing (ARR), with support for horizontal scaling
- Centralized scheduling and policy management to support backup and clone operations
- Centralized reporting, monitoring, and dashboard views
- SnapCenter 4.1 support for data protection for VMware virtual machines, SQL Server Databases, Oracle Databases, MySQL, SAP HANA, MongoDB, and Microsoft Exchange
- SnapCenter Plug-in for VMware in vCenter integration into the vSphere Web Client. All virtual machine backup and restore tasks are preformed through the web client GUI
- Using the SnapCenter Plug-in for VMware in vCenter you can:
  - Create policies, resource groups and backup schedules for virtual machines
  - Backup virtual machines, VMDKs, and datastores
  - Restore virtual machines, VMDKs, and files and folders (on Windows guest OS)
  - Attach and detach VMDK
  - Monitor and report data protection operations on virtual machines and datastores
  - Support RBAC security and centralized role delegation
  - Support guest file or folder (single or multiple) support for Windows guest OS
  - Restore an efficient storage base from primary and secondary Snapshot copies through Single File SnapRestore
  - Generate dashboard and reports that provide visibility into protected versus unprotected virtual machines and status of backup, restore, and mount jobs
  - Attach or detach virtual disks from secondary Snapshot copies
  - Attach virtual disks to an alternate virtual machine

## VMware vSphere 6.7

VMware vSphere 6.7 is the industry leading virtualization platform. VMware ESXi 6.7 is used to deploy and run the virtual machines. vCenter Server Appliance 6.7 is used to manage the ESXi hosts and virtual machines. Multiple ESXi hosts running on Cisco UCS B200 M5 blades are used to form a VMware ESXi cluster. VMware ESXi cluster pools the compute, memory and network resources from all the cluster nodes and provides a resilient platform for virtual machines running on the cluster. VMware ESXi cluster features, vSphere high availability and Distributed Resources Scheduler (DRS), contribute to the tolerance of the vSphere Cluster withstanding failures as well as distributing the resources across the VMware ESXi hosts.

## Microsoft Windows Server 2016 Hyper-V

Hyper-V is another leading virtualization platform offering from Microsoft. Hyper-V role is enabled on the Windows Server 2016 hosts to facilitate the virtual machine deployment. Multiple Windows Server 2016 hosts running on Cisco UCS B200 M5 blades are used to form Windows Server Failover Cluster (WSFC). WSFC along with Performance and Resource Management (PRO), Windows Hyper-V cluster provides high availability and resilient cluster platform for virtual machines to critical applications.

## Red Hat Enterprise Linux 7

Red Hat Enterprise Linux (RHEL) is a widely adapted Linux server operating system. In this solution, RHEL 7.4 is used as guest operating system inside the virtual machines deployed on VMware ESXi Cluster and Windows Hyper-V Cluster. Using Red Hat Enterprise Linux 7, high availability can be deployed in a variety of configurations to suit varying needs for performance, high-availability, load balancing, and file sharing. RHEL uses pacemaker, a leading high availability cluster resource manager, to manage and recover cluster resources.

## Microsoft SQL Server 2017

Microsoft SQL Server 2017 is the recent relational database engine from Microsoft. It brings in a lot of new features and enhancements to the relational and analytical engines. With SQL Server 2017, Microsoft embraces Linux operating systems and announces to support SQL Server on Linux. SQL Server 2017 is the first stable version that performs as good as it does in Windows. Currently most of the database features that are supported in Windows are also supported in Linux. This enablement on Linux removes Windows lock-in and provides flexibility to the customers to choose and deploy SQL Server databases on widely used Linux flavored operating systems.

## Solution Design

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere built on FlexPod includes NetApp All Flash FAS storage, Cisco Nexus networking, Cisco Unified Computing System, and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one datacenter rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or "flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

Figure 7 shows FlexPod components and the network connections for a configuration with the Cisco UCS 6332-16UP Fabric Interconnects. This design has end-to-end 40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnect, between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A300. This infrastructure is deployed to provide iSCSI-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 7 FlexPod With Cisco UCS 6332-16UP Fabric Interconnects




Figure 7 illustrates a base design. Each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional storage controllers or disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features.


The following components are used to the validate and test the solution:

- 1x Cisco 5108 chassis with Cisco UCS 2304 IO Modules
- 5x B200 M5 Blades with VIC 1340 and a Port Expander card
- Two Cisco Nexus 93180YC-EX switches
- Two Cisco UCS 6332-16UP fabric interconnects
- One NetApp AFF A300 (HA pair) running clustered Data ONTAP with Disk shelves and Solid-State Drives (SSD)

In this solution, both VMware ESXi and Windows Hyper-V virtual environments are tested and validated for deploying SQL Server 2017 on virtual machine running Red Hat Enterprise Linux operating system.

Figure 8 illustrates the high-level steps to deploy SQL Server on a RHEL virtual machine in a FlexPod environment with both VMware and Hyper-V hypervisors.

**Figure 8 Flowchart to Deploy SQL Server on a RHEL Virtual Machine in a FlexPod Environment**



## Deployment Hardware and Software

### Software Revisions

Table 1 lists the software revisions for this solution.

**Table 1 Software Revisions**

Layer	Device	Image	Components
Compute	Cisco UCS third-generation 6332-16UP	4.0(1c)	Includes Cisco 5108 blade chassis with UCS 2304 IO Modules  Cisco UCS B200 M5 blades with Cisco UCS VIC 1340 adapter
Network Switches	Includes Cisco Nexus 93180YC	NX-OS: 9.2.2	
Storage Controllers	NetApp AFF A300 storage controllers	Data ONTAP 9.5	
Software	Cisco UCS Manager	4.0(1c)	
	Cisco UCS Manager Plugin for VMware vSphere Web Client	2.0.4	
	VMware vSphere ESXi	6.7 U1	
	VMware vCenter	6.7 U1	
	Microsoft Windows Hyper-V	Windows Server 2016	
	NetApp Virtual Storage Console (VSC)	7.2.1	
	NetApp SnapCenter	4.1.1	
	NetApp Host Utilities Kit for RHEL 7.4 & Windows 2016	7.1	
	Red Hat Enterprise Linux 7.4	7.4	
	Microsoft SQL Server	2017	

### Configuration Guidelines

This document provides the details to configure a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either O1 or O2 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-01, VM-Host-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your

environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
[-node] <nodename>                Node
{ [-vlan-name] {<netport>|<ifgrp>} VLAN Name
| -port {<netport>|<ifgrp>}        Associated Network Port
[-vlan-id] <integer> }            Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <node01> -vlan-name i0a-<vlan id>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 3 lists the virtual machines necessary for deployment as outlined in this guide. Table 2 lists the VLANs necessary for deployment as outlined in this guide.

**Table 2 Necessary VLANs**

VLAN Name	VLAN Purpose	ID used in Validating this Document	Virtual Environment (VMware / Hyper-V)
Out of Band Mgmt	VLAN for out-of-band management interfaces	17	VMware & Hyper-V
In Band Mgmt	VLAN for in-band management interfaces	117	
Native-VLAN	VLAN to which untagged frames are assigned	2	VMware & Hyper-V
SQL-VM-MGMT	VLAN for in-band management interfaces and virtual machines interfaces	904	VMware
SQL-VM-iSCSI-A	VLAN for iSCSI A traffic for SQL virtual machines on ESXi as well as ESXi Infrastructure	3012	VMware
SQL-VM-iSCSI-B	VLAN for iSCSI B traffic for SQL virtual machines on ESXi as well as ESXi Infrastructure	3022	VMware
SQL-vMotion	VLAN for VMware vMotion	905	VMware
SQL-HV-MGMT	VLAN for in-band management interfaces and virtual machines interfaces	906	Hyper-V
SQL-HV-iSCSI-A	VLAN for iSCSI A traffic for SQL virtual machines on Hyper-V as well as Hyper-V Infrastructure	3013	Hyper-V
SQL-HV-iSCSI-B	VLAN for iSCSI B traffic for SQL virtual machines on Hyper-V as well as Hyper-V Infrastructure	3023	Hyper-V
SQL-Live-Migration	VLAN for Hyper-V virtual machine's live migration	907	Hyper-V

VLAN Name	VLAN Purpose	ID used in Validating this Document	Virtual Environment (VMware / Hyper-V)
SQL-HV-Cluster	VLAN for private cluster communication	908	Hyper-V

---

 For simplicity purposes, the same VLAN ID (SQL-VM-MGMT/SQL-HV-MGMT) is used for both in-band management interfaces of hypervisor hosts and virtual machines running the production workloads. If required, different VLANs can be used to segregate these traffics.

---

Table 3 lists the virtual machines necessary for deployment as outlined in this document.

**Table 3 Virtual Machines**

Virtual Machine Description	Host Name
Active Directory	
vCenter Server	
NetApp VSC	
NetApp SnapCenter Server	

## Physical Infrastructure

### FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain details for the prescribed and supported configuration of the NetApp AFF A300 running NetApp ONTAP 9.5.

---

 For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool \(IMT\)](#)

---

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.

---


 **Be sure to use the cabling directions in this section as a guide.**

---

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: [https://library.netapp.com/ecm/ecm\\_get\\_file/ECMM1280392](https://library.netapp.com/ecm/ecm_get_file/ECMM1280392)

The following details the cable connections used in the validation lab for the 40Gb end-to-end iSCSI topology based on the Cisco UCS 6332-16UP fabric interconnect. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlexPod infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each AFF controller has two connections to the out-of-band network switch.

Figure 9 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnects



## Network Switch Configuration

This section provides the detailed process to configure the Cisco Nexus 9000s for use in a FlexPod environment.



**Follow these steps precisely since failure to do so could result in an improper configuration.**

### Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as described in section [FlexPod Cabling](#).

### FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 9.2.2 and is valid for Cisco Nexus 93180YC-EX switches deployed with the 40Gb end-to-end topology.



**The following process includes setting up the Network Time Protocol (NTP) distribution on the in-band management VLAN. The interface-vlan feature and ntp commands are used in this setup. This procedure assumes that the default VRF is used to route the in-band management VLAN.**

### Set Up Initial Configuration

#### Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, follow these steps:

1. Configure the switch.



**On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.**

```

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>

```

```
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
2. Review the configuration summary before enabling the configuration.
Use this configuration and save it? (yes/no) [y]: Enter
```

## Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, follow these steps:

1. Configure the switch.



**On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.**

---

```
Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
```

```
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L3]: L2
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
2. Review the configuration summary before enabling the configuration.
Use this configuration and save it? (yes/no) [y]: Enter
```

## FlexPod Cisco Nexus Switch Configuration

### Enable Licenses

#### Cisco Nexus A and Cisco Nexus B

To license the Cisco Nexus switches, follow these steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature interface-vlan
feature lacp
feature vpc
feature lldp
```

### Set Global Configurations

#### Cisco Nexus A and Cisco Nexus B

To set global configurations, follow this step on both switches:

1. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

## Create VLANs

### Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
vlan <Native-VLAN-id>
name Native-VLAN

vlan <SQL-VM-MGMT-id>
name SQL-VM-MGMT

vlan <SQL-vMotion-id>
name SQL-vMotion

vlan <SQL-VM-iSCSI-A-id>
name SQL-VM-iSCSI-A

vlan <SQL-VM-iSCSI-B-id>
name SQL-VM-iSCSI-B

vlan <SQL-HV-MGMT-id>
name SQL-HV-MGMT

vlan <SQL-Live-Migration-id>
name SQL-Live-Migration

vlan <SQL-HV-iSCSI-A-id>
name SQL-HV-iSCSI-A

vlan <SQL-HV-iSCSI-B-id>
```

```
name SQL-HV-iSCSI-B
vlan <SQL-HV-Cluster-id>
name SQL-HV-Cluster
exit
```

## Add NTP Distribution Interface

### Cisco Nexus A

From the global configuration mode, run the following commands:

```
ntp source <switch-a-ntp-ip>
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
```

### Cisco Nexus B

From the global configuration mode, run the following commands:

```
ntp source <switch-b-ntp-ip>
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
```

## Add Individual Port Descriptions for Troubleshooting

### Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, follow this step:



**In this step and in the later sections, configure the AFF nodename <st-node> and Cisco UCS 6332-16UP fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.**

---

1. From the global configuration mode, run the following commands:

```
interface Eth1/51
description <st-node>-1:e2a
interface Eth1/52
description <st-node>-2:e2a
interface Eth1/53
description <ucs-clustername>-A:1/39
```

```
interface Eth1/54
description <ucs-clustername>-B:1/39
interface Eth1/43
description <nexus-hostname>-B:1/43
interface Eth1/44
description <nexus-hostname>-B:1/44
interface Eth1/45
description <nexus-hostname>-B:1/45
interface Eth1/46
description <nexus-hostname>-B:1/46
exit
```

### Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B, follow this step:

1. From the global configuration mode, run the following commands:

```
interface Eth1/51
description <st-node>-1:e2e
interface Eth1/52
description <st-node>-2:e2e
interface Eth1/53
description <ucs-clustername>-A:1/40
interface Eth1/54
description <ucs-clustername>-B:1/40
interface Eth1/43
description <nexus-hostname>-A:1/43
interface Eth1/44
description <nexus-hostname>-A:1/44
interface Eth1/45
description <nexus-hostname>-A:1/45
interface Eth1/46
description <nexus-hostname>-A:1/46
exit
```

## Create Port Channels

### Cisco Nexus A and Cisco Nexus B

To create the necessary port channels between devices, follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/43-46
channel-group 10 mode active
no shutdown
```

```
interface Po151
description <st-node>-1
interface Eth1/51
channel-group 151 mode active
no shutdown
```

```
interface Po152
description <st-node>-2
interface Eth1/52
channel-group 152 mode active
no shutdown
```

```
interface Po153
description <ucs-clustername>-A
interface Eth1/53
channel-group 153 mode active
no shutdown
```

```
interface Po154
description <ucs-clustername>-B
interface Eth1/54
channel-group 154 mode active
no shutdown
```

```
exit
copy run start
```

## Configure Port Channel Parameters

### Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <SQL-VM-MGMT-id>,<SQL-VM-iSCSI-A-id>,<SQL-VM-iSCSI-B-
id>,<SQL-vMotion-id>,<SQL-HV-MGMT-id>,<SQL-HV-iSCSI-A-id>,<SQL-HV-iSCSI-B-id>,<SQL-
Live-Migration-id>,<SQL-HV-Cluster-id>
spanning-tree port type network
```

```
interface Po151
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <SQL-VM-iSCSI-A-id>,<SQL-VM-iSCSI-B-id>,<SQL-HV-
iSCSI-A-id>,<SQL-HV-iSCSI-B-id>
spanning-tree port type edge trunk
mtu 9216
```

```
interface Po152
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <SQL-VM-iSCSI-A-id>,<SQL-VM-iSCSI-B-id>,<SQL-HV-
iSCSI-A-id>,<SQL-HV-iSCSI-B-id>
spanning-tree port type edge trunk
mtu 9216
```

```
interface Po153
switchport mode trunk
switchport trunk native vlan 2
```

```
switchport trunk allowed vlan <SQL-VM-MGMT-id>, <SQL-VM-iSCSI-A-id>, <SQL-VM-iSCSI-B-
id>, <SQL-vMotion-id>, <SQL-HV-MGMT-id>, <SQL-HV-iSCSI-A-id>, <SQL-HV-iSCSI-B-id>, <SQL-
Live-Migration-id>, <SQL-HV-Cluster-id>

spanning-tree port type edge trunk

mtu 9216

interface Po154
switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <SQL-VM-MGMT-id>, <SQL-VM-iSCSI-A-id>, <SQL-VM-iSCSI-B-
id>, <SQL-vMotion-id>, <SQL-HV-MGMT-id>, <SQL-HV-iSCSI-A-id>, <SQL-HV-iSCSI-B-id>, <SQL-
Live-Migration-id>, <SQL-HV-Cluster-id>

spanning-tree port type edge trunk

mtu 9216

exit

copy run start
```

## Configure Virtual Port Channels

### Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, follow this step:

1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po151
vpc 151
interface Po152
vpc 152
```

```
interface Po153
vpc 153
interface Po154
vpc 154
exit
copy run start
```

### Cisco Nexus B

To configure vPCs for switch B, follow this step:

1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po151
vpc 151
interface Po152
vpc 152
interface Po153
vpc 153
interface Po154
vpc 154
exit
copy run start
```

### Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is completed.

# Storage Configuration

---

## NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.



**Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the [HWU application](#) at the [NetApp Support](#) site.**

---

1. Access the [HWU](#) application to view the system configuration guides. Click the Platforms tab to view the compatibility between different version of the ONTAP software and the NetApp storage Platforms with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

## Controllers and Disk Shelves

Follow the physical installation procedures for the controllers found in the [AFF A300 Series product documentation](#) in the [NetApp AFF and FAS Documentation Center](#).

### AFF A300 Controllers

See the Installation and Setup section in the [NetApp AFF and FAS System Documentation Center](#) for planning the physical location of the storage systems.

### Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by the AFF A300 is available at the [NetApp Support](#) site.

When using SAS disk shelves with NetApp storage controllers, refer to the [SAS Disk Shelves Universal SAS and ACP Cabling Guide](#) for proper cabling guidelines.

## Clustered Data ONTAP 9.5

### Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the [ONTAP 9.5 Software Setup Guide](#) in the [ONTAP 9 Documentation Center](#). You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

### Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [ONTAP 9.5 Software Setup Guide](#) to learn about configuring ONTAP. Table 4 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

**Table 4 ONTAP Software Installation Prerequisites**

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Data ONTAP 9.5 URL	<url-boot-software>



Pursuant to best practices, NetApp recommends the following command on the LOADER prompt of the NetApp controllers to assist with LUN stability during copy operations. To access the LOADER prompt, connect to the controller through serial console port or Service Processor connection and press Ctrl-C to halt the boot process when prompted: `setenv bootarg.tmgr.disable_pit_hp 1`



For more information about the workaround, please see the NetApp public report. Note that a NetApp login is required to view the report: <http://nt-ap.com/2w6myr4>



For more information about Windows Offloaded Data Transfers see: [https://technet.microsoft.com/en-us/library/hh831628\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831628(v=ws.11).aspx)

### Configure Node 01

To configure node 01, follow these steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when the following message displays:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.5 is the version being booted, select option 8 and `y` to reboot the node; then continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select `e0M` for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for `e0M`.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

9. Enter the URL where the software can be found.



**This web server must be pingable.**

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.



**When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.**

13. Press Ctrl-C when the following message displays:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.



**The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node 02 configuration while the disks for node 01 are zeroing.**

## Configure Node 02

To configure node 02, follow these steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when the following message displays:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



**If ONTAP 9.5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.5 is the version being booted, select option 8 and `y` to reboot the node. Then continue with step 14.**

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M.

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

9. Enter the URL where the software can be found.



**This web server must be pingable.**

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.



**When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.**

13. Press Ctrl-C when you see this message:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.



**The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.**

## Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.5 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```

Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem
occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created

Use your web browser to complete cluster setup by accessing https://<node01-mgmt-ip>

Otherwise press Enter to complete cluster setup using the command line interface:

```

2. To complete the cluster setup, open a web browser and navigate to <https://<node01-mgmt-ip>>.

**Table 5 Cluster Create in ONTAP Prerequisites**

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>

Cluster Detail	Cluster Detail Value
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-SP-ip>
Node 02 service processor IP address	<node02-SP-ip>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server IP address	<ntp-ip>



Cluster setup can also be done by using the CLI. This document describes the cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup on the Welcome screen.

Figure 10 Setup NetApp Storage Cluster-1

Cluster Setup Workflow - X

Not Secure | https://192.168.156.61/sysmgr/SysMgr.html

NetApp OnCommand System Manager

Getting Started

Language English (English)

### Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:

- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications

For information related to setting up the cluster, [click here](#)

#### Template File

Browse to select a .csv file...

To download the template, click [file.csv](#) or [file.xlsx](#)

Important: You can download the template in ".csv" or ".xlsx" format. However, you can upload only those templates that are in ".csv" format.

Guided Setup

Click to set up the cluster

4. On the Cluster page, do as follows:
  - a. Enter the cluster and node names.
  - b. Select the cluster configuration.
  - c. Enter and confirm the password.
  - d. (Optional) Enter the cluster base and feature licenses.

Figure 11 Setup NetApp Storage Cluster-2



Cluster Name

Nodes

*i* Not sure all nodes have been discovered? [Refresh](#)

AFF-A300	721653000058	HA PAIR	AFF-A300	721653000057
<input checked="" type="checkbox"/>	<input type="text" value="bb04-aff300-1"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="text" value="bb04-aff300-2"/>

Cluster Configuration:  Switched Cluster  Switchless Cluster

*i* Ensure that the hardware connectivity is set up for the two-node switchless cluster.

*?* Username

Password

Confirm Password

Cluster Base License (Optional)

*i* For any queries related to licenses, contact [mysupport.netapp.com](https://mysupport.netapp.com)

Feature Licenses (Optional)

*i* Cluster Base License is mandatory to add Feature Licenses.

---



The nodes are discovered automatically; if they are not, click the Refresh link. By default, the cluster interfaces will be created on all the new factory shipping storage controllers.

---



If all the nodes are not discovered, then configure the cluster by using the CLI.

---




Cluster license and feature licenses can also be installed after completing the cluster creation.

---

5. Click Submit.
6. On the Network page, complete the following sections:
  - Cluster Management: Enter the IP address, netmask, gateway and port details.
  - Node Management: Enter the node management IP addresses and port details for all the nodes.
  - Service Processor Management: Enter the IP addresses for all the nodes.
  - DNS Details: Enter the DNS domain names and server address.
  - NTP Details: Enter the primary and alternate NTP server.
7. Click Submit.

**Figure 12 Setup NetApp Storage Cluster-3**  
**Guided Setup to Configure a Cluster**

Provide the information required below to configure your cluster:



**Network (Management)**

**IP Addresses (IPv4) required** Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

**IP Address Range**  You must enter the default network details manually.

	IP Address	Netmask	Gateway (Optional)	Port
<b>Cluster Management</b>	192.168.156.60	255.255.255.0	192.168.156.1	e0c
⚠ Ensure that the cluster management LIF is reachable or a Gateway is configured for the same subnet in which the cluster management LIF is present.				
<b>Node Management</b>	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
bb04-aff300-1	192.168.156.61	e0M		
bb04-aff300-2	192.168.156.62	e0M		
<b>Service Processor Management</b>	Default values have been detected for the Service Processor.			
	<input type="checkbox"/> Override the default values (Gateway is mandatory)			
	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
bb04-aff300-1	192.168.156.58			
bb04-aff300-2	192.168.156.59			

**DNS Details**

DNS Domain Names

DNS Server IP Address

**NTP Details**

Primary NTP Server

Alternative NTP Server (Optional)

8. On the Support page, configure the AutoSupport and Event Notifications sections.

Figure 13 Setup NetApp Storage Cluster-4

### Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



#### ? AutoSupport

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

#### ? Event Notifications

Notify me through:

<input checked="" type="checkbox"/>	<b>Email</b>	<b>SMTP Mail Host</b> <input type="text" value="testvikings.smtp.cisco.com"/>	<b>Email Addresses</b> <input type="text" value="adminvikings@cisco.com"/>
-------------------------------------	--------------	--	---

<input type="checkbox"/>	<b>SNMP</b>	<b>SNMP Trap Host</b> <input type="text"/>
--------------------------	-------------	---

<input type="checkbox"/>	<b>Syslog</b>	<b>Syslog Server</b> <input type="text"/>
--------------------------	---------------	--

**Submit**

9. Click Submit.
10. On the Summary page, review the configuration details if needed.

Figure 14 Setup NetApp Storage Cluster-5

## Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



[Click here to view the summary](#)

The next step will be to configure your aggregates, SVM and Storage Objects.  
Click the button below to start provisioning your storage.

[Manage your cluster](#)



The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

### Log into the Cluster

To log in to the cluster, follow these steps:

1. Open an SSH connection to either the cluster IP or host name.
2. Log in to the admin user with the password you provided earlier.

### Zero All Spare Disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```



Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an AFF configuration. Disk autoassign should have assigned one data partition to each node in an HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

### Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard unified target adapter 2 (UTA2), follow these steps:

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin show` command:

```
ucadmin show
Node          Adapter      Current      Current      Pending      Pending      Admin
              Mode         Type         Mode         Type         Status
```

```

-----
<st-node01> 0e cna target - - online
<st-node01> 0f cna target - - online
<st-node01> 0g cna target - - online
<st-node01> 0h cna target - - online
<st-node02> 0e cna target - - online
<st-node02> 0f cna target - - online
<st-node02> 0g cna target - - online
<st-node02> 0h cna target - - online
8 entries were displayed.

```

- Verify that the Current Mode and Current Type properties for all ports are set properly. Set the ports used for iSCSI connectivity to mode `cna`. The port type for all protocols should be set to `target`. Change the port personality by running the following command:

```
ucadmin modify -node <home-node-of-the-port> -adapter <port-name> -mode {fc|cna} -type target
```



**The ports must be offline to run this command. To take an adapter offline, run the `fc` adapter `modify -node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, `0e` and `0f`). After conversion, a reboot is required, and the ports must be brought back to the up state.**

## Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, follow this step:



**An SVM is referred to as a Vserver (or `vserver`) in the GUI and CLI.**

Run the following command:

```
network interface modify -vserver <clustername> -lif cluster_mgmt -auto-revert true
```

## Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, `e0d`, `e2a`, and `e2e`) should be removed from the default broadcast domain, leaving just the management network ports (`e0c` and `e0M`). To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports <st-node01>:e0d,<st-node01>:e0e,<st-
node01>:e0f,<st-node01>:e0g,<st-node01>:e0h,<st-node01>:e2a, <st-node01>:e2e,<st-node02>:e0d,<st-
node02>:e0e,<st-node02>:e0f,<st-node02>:e0g,<st-node02>:e0h <st-node02>:e2a,<st-node02>:e2e
broadcast-domain show
```

## Set Up Service Processor Network Interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <st-node01> -address-family IPv4 -enable true -dhcp
none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>

system service-processor network modify -node <st-node02> -address-family IPv4 -enable true -dhcp
none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

## Create Aggregates

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain.

To create new aggregates, follow these steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <st-node01> -diskcount <num-disks>
aggr create -aggregate aggr1_node02 -node <st-node02> -diskcount <num-disks>
```



You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate.



For all flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions. For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type.



Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all but one remaining disk (spare) assigned to the controller.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

2. (Optional) Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02. The aggregate will be automatically renamed if System guided setup is used.

```
aggr show
aggr rename -aggregate aggr0 -newname <node01-rootaggrname>
```

## Verify Storage Failover

To confirm that storage failover is enabled, run the following commands for a failover pair:

1. Verify the status of the storage failover.

```
storage failover show
```



Both <st-node01> and <st-node02> must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <st-node01> -enabled true
```



Enabling failover on one node enables it for both nodes.

3. Verify the HA status for a two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.
5. Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <node02-mgmt-ip> -node <st-node01>
storage failover modify -hwassist-partner-ip <node01-mgmt-ip> -node <st-node02>
```

## Disable Flow Control on 10GbE and 40GbE Ports

NetApp recommends disabling flow control on all the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, follow these steps:

1. Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

2. Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show -fields flowcontrol-admin
```

## Disable Unused FCoE Capability on CNA Ports

If the UTA2 port is set to CNA mode and is only expected to handle Ethernet data traffic (for example NFS), then the unused FCoE capability of the port should be disabled by setting the corresponding FCP adapter to state down with the `fc adapter modify` command. Here are some examples:

```
fc adapter modify -node <st-node01> -adapter 0e -status-admin down
fc adapter modify -node <st-node01> -adapter 0f -status-admin down
fc adapter modify -node <st-node01> -adapter 0g -status-admin down
```

```

fcpx adapter modify -node <st-node01> -adapter 0h -status-admin down
fcpx adapter modify -node <st-node02> -adapter 0e -status-admin down
fcpx adapter modify -node <st-node02> -adapter 0f -status-admin down
fcpx adapter modify -node <st-node02> -adapter 0g -status-admin down
fcpx adapter modify -node <st-node02> -adapter 0h -status-admin down
fcpx adapter show -fields -status-admin

```

## Configure Network Time Protocol

To configure time synchronization on the cluster, follow these steps:

1. Set the time zone for the cluster.

```
timezone <timezone>
```



**For example, in the eastern United States, the time zone is `America/New_York`.**

2. Set the date for the cluster.

```
date <ccyymmddhhmm.ss>
```



**The format for the date is `<[Century] [Year] [Month] [Day] [Hour] [Minute] . [Second]>` (for example, `201309081735.17`).**

3. Configure the NTP servers for the cluster.

```

cluster time-service ntp server create -server <switch-a-ntp-ip>
cluster time-service ntp server create -server <switch-b-ntp-ip>

```

## Configure Simple Network Management Protocol

To configure the SNMP, follow these steps:

1. Configure basic SNMP information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```

snmp contact <snmp-contact>
snmp location "<snmp-location>"
snmp init 1
options snmp.enable on

```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <oncommand-um-server-fqdn>
```

## Configure SNMPv1 Access

To configure SNMPv1 access, set the shared, secret plain-text password (called a community).

```
snmp community add ro <snmp-community>
```

## Configure AutoSupport

NetApp AutoSupport sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <mailhost> -transport https -support
enable -noteto <storage-admin-email>
```

## Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command to enable CDP on ONTAP:

```
node run -node * options cdpd.enable on
```



**To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.**

## Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands to create a broadcast domain for NFS and iSCSI on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

## Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2e

ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2e

ifgrp show
```

## Create VLANs

To create VLANs, create NFS VLAN ports and add them to the NFS broadcast domain:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000
network port vlan create -node <st-node01> -vlan-name a0a-<infra-nfs-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-nfs-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <st-node01>:a0a-<infra-nfs-vlan-id>,
<st-node02>:a0a-<infra-nfs-vlan-id>
```

To create VLANs, create iSCSI VLAN ports and add them to the iSCSI broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<SQL-VM-iSCSI-A-id>
network port vlan create -node <st-node02> -vlan-name a0a-<SQL-VM-iSCSI-A-id>
network port vlan create -node <st-node01> -vlan-name a0a-<SQL-VM-iSCSI-B-id>
network port vlan create -node <st-node02> -vlan-name a0a-<SQL-VM-iSCSI-B-id>
network port vlan create -node <st-node01> -vlan-name a0a-<SQL-HV-iSCSI-A-id>
network port vlan create -node <st-node02> -vlan-name a0a-<SQL-HV-iSCSI-A-id>
network port vlan create -node <st-node01> -vlan-name a0a-<SQL-HV-iSCSI-B-id>
network port vlan create -node <st-node02> -vlan-name a0a-<SQL-HV-iSCSI-B-id>

broadcast-domain add-ports -broadcast-domain SQL-VM-iSCSI-A -ports <st-node01>:a0a-<SQL-VM-iSCSI-A-
id>, <st-node02>:a0a-<SQL-VM-iSCSI-A-id>
broadcast-domain add-ports -broadcast-domain SQL-VM-iSCSI-B -ports <st-node01>:a0a-<SQL-VM-iSCSI-B-
id>, <st-node02>:a0a-<SQL-VM-iSCSI-B-id>
broadcast-domain add-ports -broadcast-domain SQL-HV-iSCSI-A -ports <st-node01>:a0a-<SQL-HV-iSCSI-A-
id>, <st-node02>:a0a-<SQL-HV-iSCSI-A-id>
```

```

broadcast-domain add-ports -broadcast-domain SQL-HV-iSCSI-B -ports <st-node01>:a0a-<SQL-HV-iSCSI-B-id>, <st-node02>:a0a-<SQL-HV-iSCSI-B-id>
broadcast-domain add-ports -broadcast-domain SQL-VM-MGMT -ports <st-node01>:a0a-<SQL-VM-MGMT-id>, <st-node02>:a0a-<SQL-VM-MGMT-id>
broadcast-domain add-ports -broadcast-domain SQL-HV-MGMT -ports <st-node01>:a0a-<SQL-HV-MGMT-id>, <st-node02>:a0a-<SQL-HV-MGMT-id>
broadcast-domain add-ports -broadcast-domain Infra-NFS -ports <st-node01>:a0a-<Infra-NFS-id>, <st-node02>:a0a-<Infra-NFS-id>

```

## Infrastructure Storage Virtual Machine (Optional)



Create a separate SVM for hosting infrastructure services (virtual machines) on dedicated UCS servers. If you are planning to combine infrastructure service virtual machines along with workload virtual machines, this step can be skipped.

### Create Storage Virtual Machine

To create an infrastructure SVM, follow these steps:

1. Run the `vserver create` command.

```

vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-security-style unix

```

2. Select the SVM data protocols to configure, keeping iSCSI and NFS.

```

vserver remove-protocols -vserver Infra-SVM -protocols fcp,cifs,ndmp

```

3. Add the two data aggregates to the Infra-SVM aggregate list for the NetApp VSC.

```

vserver modify -vserver Infra-SVM -aggr-list aggr1_node01,aggr1_node02

```

4. Enable and run the NFS protocol in the Infra-SVM.

```

nfs create -vserver Infra-SVM -udp disabled

```

5. Turn on the SVM vStorage parameter for the NetApp NFS vStorage APIs – Array integration VAAI plug-in.

```

vserver nfs modify -vserver Infra-SVM -vstorage enabled
vserver nfs show

```

### Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```

volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP

```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```

job schedule interval create -name 15min -minutes 15

```

3. Create the mirroring relationships.

```

snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m01 -type LS -
schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m02 -type LS -
schedule 15min

```

4. Initialize the mirroring relationship.

```

snapmirror initialize-ls-set -source-path Infra-SVM:rootvol
snapmirror show

```

## Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

```

iscsi create -vserver Infra-SVM
iscsi show

```

## Configure HTTPS Access

To configure secure access to the storage controller, follow these steps:

1. Increase the privilege level to access the certificate commands.

```

set -privilege diag
Do you want to continue? {y|n}: y

```

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <serial-number>) by running the following command:

```

security certificate show

```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```

security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -type server -
serial <serial-number>

```



**Deleting expired certificates before creating new certificates is a best practice. Run the `security certificate delete` command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.**

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```

security certificate create -common-name <cert-common-name> -type server -size
2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -
organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days
<cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-SVM

```

5. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the `security certificate show` command.

6. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

7. Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```



**It is normal for some of these commands to return an error message stating that the entry does not exist.**

8. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
```

## Configure NFSv3

To configure NFSv3 on the SVM, follow these steps:

1. Create a new rule for the infrastructure NFS subnet in the default export policy.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol nfs -clientmatch <infra-nfs-subnet-cidr> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

2. Assign the FlexPod export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```

## Create FlexVol Volumes

The following information is required to create a NetApp FlexVol volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

To create a FlexVol volume, run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -size 500GB -state online -policy default -junction-path /infra_datastore_1 -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume infra_datastore_2 -aggregate aggr1_node02 -size 500GB -state online -policy default -junction-path /infra_datastore_2 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size 100GB -state online -policy default -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path Infra-SVM:rootvol
```

## Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB -ostype vmware -
space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-02 -size 15GB -ostype vmware -
space-reserve disabled
```

## Schedule Deduplication

On NetApp AFF systems, deduplication is enabled by default. To schedule deduplication, follow these steps:

1. After the volumes are created, assign a once-a-day deduplication schedule to `esxi_boot` and `infra_datastore_1`:

```
efficiency modify -vserver Infra-SVM -volume esxi_boot -schedule sun-sat@0
efficiency modify -vserver Infra-SVM -volume infra_datastore_1 -schedule sun-sat@0
```

## Create Management LIF

To create SVM management LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif IB-MGMT -role data -data-protocol none -home-node
<st-node01> -home-port a0a-<IB-MGMT-vlan-id> -address <Infra-SVM-Mgmt_ip> -netmask <Infra-SVM-
Mgmt_mask> -status-admin up -failover-policy system-defined -firewall-policy mgmt -auto-revert true
```

## Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<infra-iscsi-a-vlan-id> -address <var_node01_infra-iscsi_lif01a_ip> -
netmask <var_node01_iscsi_lif01a_mask> -status-admin up -failover-policy disabled -firewall-policy
data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-
node <st-node01> -home-port a0a-<infra-iscsi-b-vlan-id> -address <var_node01_infra-iscsi_lif01b_ip> -
netmask <var_node01_iscsi_lif01b_mask> -status-admin up -failover-policy disabled -firewall-policy
data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -address <var_node02_infra-iscsi_lif02a_ip> -
netmask <var_node02_iscsi_lif02a_mask> -status-admin up -failover-policy disabled -firewall-policy
data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-
node <st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -address <var_node02_infra-iscsi_lif02b_ip> -
netmask <var_node02_iscsi_lif02b_mask> -status-admin up -failover-policy disabled -firewall-policy
data -auto-revert false

network interface show
```

## Create NFS LIF

To create an NFS LIF, run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_lif01 -role data -data-protocol nfs -home-node
<st-node01> -home-port a0a-<infra-nfs-vlan-id> -address <node01-nfs_lif01-ip> -netmask <node01-
nfs_lif01-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true
```

```
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data -data-protocol nfs -home-node
<st-node02> -home-port a0a-<infra-nfs-vlan-id> -address <node02-nfs_lif02-ip> -netmask <node02-
nfs_lif02-mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-
revert true

network interface show
```

## Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif svm-mgmt -role data -data-protocol none -home-node
<st-node02> -home-port e0c -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up -
failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```



**The SVM management IP in this step should be in the same subnet as the storage cluster management IP.**

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>

network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver Infra-SVM
```



**A cluster serves data through at least one and possibly several SVMs. We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.**

## Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into following tables.

**Table 6 iSCSI IQN and iSCSI LIFs**

IQN or LIF IP	Info
Infra-SVM IQN	
infra-iscsi_lif01a_ip	
infra-iscsi_lif02a_ip	
infra-iscsi_lif01b_ip	

IQN or LIF IP	Info
infra-iscsi_lif02b_ip	



To obtain the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

**Table 7 vNIC iSCSI IQNs for Fabric A and Fabric B**

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-Infra-01		<vm-host-infra-01-iqn>
VM-Host-Infra-02		<vm-host-infra-02-iqn>



To obtain the iSCSI vNIC IQN information in the Cisco UCS Manager GUI, go to **Servers > Service Profiles > root**. Click each service profile and then select the “iSCSI vNICs” tab. The initiator name is displayed at the top of the page under the service profile initiator name.

## SVM for SQL Linux Virtual Machines on ESXi

### Create Storage Virtual Machines

To create an infrastructure SVM, follow these steps:

1. Run the `vserver create` command.

```
vserver create -vserver ESXi-Work-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-
security-style unix
```

2. Select the SVM data protocols to configure, keeping iSCSI only.

```
vserver remove-protocols -vserver ESXi-Work-SVM -protocols nfs,fc,fcifs,ndmp
```

3. Add the two data aggregates to the ESXi-Work-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver ESXi-Work-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the ESXi-Work-SVM.

```
nfs create -vserver ESXi-Work-SVM -udp disabled
```

### Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver ESXi-Work-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver ESXi-Work-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path ESXi-Work-SVM:rootvol -destination-path ESXi-Work-SVM:rootvol_m01 -
type LS -schedule 15min
snapmirror create -source-path ESXi-Work-SVM:rootvol -destination-path ESXi-Work-SVM:rootvol_m02 -
type LS -schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path ESXi-Work-SVM:rootvol
snapmirror show
```

## Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the IQN for the SVM.

```
iscsi create -vserver ESXi-Work-SVM
iscsi show
```

## Configure HTTPS Access

To configure secure access to the storage controller, follow these steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <serial-number>) by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver ESXi-Work-SVM -common-name ESXi-Work-SVM -ca ESXi-Work-SVM -type
server -serial <serial-number>
```



**Deleting expired certificates before creating new certificates is a best practice. Run the `security certificate delete` command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.**

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the ESXi-Work-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-
country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -
email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver ESXi-
Work-SVM
```

5. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the `security certificate show` command.

6. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -
serial <cert-serial> -common-name <cert-common-name>
```

7. Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```



**It is normal for some of these commands to return an error message stating that the entry does not exist.**

8. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
```

## Create FlexVol Volumes

The following information is required to create a NetApp FlexVol volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

To create a FlexVol volume, run the following commands:

```
volume create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -aggregate aggr1_node02 -size 1TB -state
online -policy default -junction-path /esxi_hosts_vms -space-guarantee none -percent-snapshot-space 0

volume create -vserver ESXi-Work-SVM -volume esxi_swap -aggregate aggr1_node01 -size 750GB -state
online -policy default -junction-path /esxi_swap -space-guarantee none -percent-snapshot-space 0

volume create -vserver ESXi-Work-SVM -volume sql_esxi_data -aggregate aggr1_node02 -size 2.5TB -state
online -policy default -junction-path /sql_esxi_data -space-guarantee none -percent-snapshot-space 0

volume create -vserver ESXi-Work-SVM -volume sql_esxi_log -aggregate aggr1_node02 -size 1.5TB -state
online -policy default -junction-path /sql_esxi_log -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path ESXi-Work-SVM:rootvol
```

## Create ESXi Boot LUNs

To create three boot LUNs for ESXi hosts, run the following commands:

```
lun create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host1 -size 25GB -ostype vmware -
space-reserve disable
lun create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host2 -size 25GB -ostype vmware -
space-reserve disable
lun create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host3 -size 25GB -ostype vmware -
space-reserve disable
```

## Create Storage Volume for SQL Virtual Machine Datastore and Swap LUNs

To create a SQL Virtual Machine Datastore LUN and Swap LUN, run the following commands:

```
lun create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun sql-vms-datastore -size 500GB -ostype
vmware -space-reserve disable
lun create -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-vms-swap -size 500GB -ostype
vmware -space-reserve disable
```

## Create SQL Data and Log LUNs

To create a SQL Virtual Machine Datastore LUN and Swap LUN, run the following commands:

```
lun create -vserver ESXi-Work-SVM -volume sql_esxi_data -lun SQL-ESXi-Data-01 -size 400GB -ostype
linux -space-reserve diable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_data -lun SQL-ESXi-Data-02 -size 400GB -ostype
linux -space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_data -lun SQL-ESXi-Data-03 -size 400GB -ostype
linux -space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_data -lun SQL-ESXi-Data-04 -size 400GB -ostype
linux -space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_log -lun SQL-ESXi-Log-01 -size 150GB -ostype linux
-space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_log -lun SQL-ESXi-Log-02 -size 150GB -ostype linux
-space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_log -lun SQL-ESXi-Log-03 -size 150GB -ostype linux
-space-reserve disable
lun create -vserver ESXi-Work-SVM -volume sql_esxi_log -lun SQL-ESXi-Log-04 -size 150GB -ostype linux
-space-reserve disable
```

## Schedule Deduplication

On NetApp AFF systems, deduplication is enabled by default. To schedule deduplication, follow these steps:

1. After the volumes are created, assign a once-a-day deduplication schedule to `esxi_boot` and `infra_datastore_1`:

```
efficiency modify -vserver ESXi-Work-SVM -volume esxi_hosts_vms -schedule sun-sat@0
efficiency modify -vserver ESXi-Work-SVM -volume esxi_swap -schedule sun-sat@0
efficiency modify -vserver ESXi-Work-SVM -volume sql_esxi_data -schedule sun-sat@0
efficiency modify -vserver ESXi-Work-SVM -volume sql_esxi_log -schedule sun-sat@0
```

## Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver ESXi-Work-SVM -lif esxi-iscsi-lif01a -role data -data-protocol
iscsi -home-node <st-node01> -home-port a0a-<SQL-VM-iSCSI-A-id> -address <var_node01_esxi-iscsi-
lif01a_ip> -netmask <var_node01_esxi-iscsi_lif01a_mask> -status-admin up -failover-policy disabled -
firewall-policy data -auto-revert false

network interface create -vserver ESXi-Work-SVM -lif esxi-iscsi_lif01b -role data -data-protocol
iscsi -home-node <st-node01> -home-port a0a-<SQL-VM-iSCSI-B-id> -address <var_node01_esxi-
iscsi_lif01b_ip> -netmask <var_node01_esxi-iscsi_lif01b_mask> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false

network interface create -vserver ESXi-Work-SVM -lif esxi-iscsi_lif02a -role data -data-protocol
iscsi -home-node <st-node02> -home-port a0a-<SQL-VM-iSCSI-A-id> -address <var_node02_esxi-
iscsi_lif02a_ip> -netmask <var_node02_esxi-iscsi_lif02a_mask> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false

network interface create -vserver ESXi-Work-SVM -lif esxi-iscsi_lif02b -role data -data-protocol
iscsi -home-node <st-node02> -home-port a0a-<SQL-VM-iSCSI-B-id> -address <var_node02_esxi-
iscsi_lif02b_ip> -netmask <var_node02_esxi-iscsi_lif02b_mask> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false

network interface show
```

## Add SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver ESXi-Work-SVM -lif ESXi-Work-SVM-mgmt -role data -data-protocol
none -home-node <st-node01> -home-port e0c -address <ESXi-Work-SVM-Mgmt_ip> -netmask <ESXi-Work-SVM-
Mgmt_mask> -status-admin up -failover-policy system-defined -firewall-policy mgmt -auto-revert true
```



The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver ESXi-Work-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver ESXi-Work-SVM
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver ESXi-Work-SVM
```



A cluster serves data through at least one and possibly several SVMs. We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.

## Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into following tables.

**Table 8 iSCSI IQN and iSCSI LIFs**

IQN or LIF IP	Info
ESXi-Work-SVM IQN	
esxi-iscsi_lif01a_ip	
esxi-iscsi_lif02a_ip	
esxi-iscsi_lif01b_ip	
esxi-iscsi_lif02b_ip	



To obtain the iSCSI IQN, run `iscsi show` command on the storage cluster management interface.

**Table 9 vNIC iSCSI IQNs for Fabric A and Fabric B**

Cisco UCS Service Profile Name	iSCSI IQN	Variables
esxi-host1		<esxi-host1-iqn>
esxi-host2		<esxi-host2-iqn>
esxi-host3		<esxi-host3-iqn>



To obtain the iSCSI vNIC IQN information in the Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then select the iSCSI vNICs tab. The initiator name is displayed at the top of the page under the service profile initiator name.

## SVM for SQL Linux Virtual Machines on Hyper-V

### Create Storage Virtual Machine

To create an infrastructure SVM, follow these steps:

1. Run the `vserver create` command.

```
vserver create -vserver Hyperv-Work-SVM -rootvolume rootvol -aggregate aggr1_node01 -rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping iSCSI only.

```
vserver remove-protocols -vserver Hyperv-Work-SVM -protocols nfs,fc,fcifs,ndmp
```

3. Add the two data aggregates to the Hyperv-Work-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Hyperv-Work-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the Hyperv-Work-SVM.

```
nfs create -vserver Hyperv-Work-SVM -udp disabled
```

### Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, follow these steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Hyperv-Work-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Hyperv-Work-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Hyperv-Work-SVM:rootvol -destination-path Hyperv-Work-SVM:rootvol_m01
-type LS -schedule 15min
snapmirror create -source-path Hyperv-Work-SVM:rootvol -destination-path Hyperv-Work-SVM:rootvol_m02
-type LS -schedule 15min
```

- Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Hyperv-Work-SVM:rootvol
snapmirror show
```

## Create Block Protocol (iSCSI) Service

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the IQN for the SVM.

```
iscsi create -vserver Hyperv-Work-SVM
iscsi show
```

## Configure HTTPS Access

To configure secure access to the storage controller, follow these steps:

- Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

- Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <serial-number>) by running the following command:

```
security certificate show
```

- For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Hyperv-Work-SVM -common-name Hyperv-Work-SVM -ca Hyperv-Work-SVM
-type server -serial <serial-number>
```



**Deleting expired certificates before creating new certificates is a best practice. Run the `security certificate delete` command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.**

- To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Hyperv-Work-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country> -state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Hyperv-Work-SVM
```

- To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the `security certificate show` command.
- Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

7. Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```



**It is normal for some of these commands to return an error message stating that the entry does not exist.**

8. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
```

## Create FlexVol Volumes

The following information is required to create a NetApp FlexVol volume:

- The volume name
- The volume size
- The aggregate on which the volume exists

To create a FlexVol volume, run the following commands:

```
volume create -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -aggregate aggr1_node02 -size 3TB -
state online -policy default -junction-path /hyperv_hosts_vms -space-guarantee none -percent-
snapshot-space 0

volume create -vserver Hyperv-Work-SVM -volume sql_hyperv_data -aggregate aggr1_node02 -size 2.5TB -
state online -policy default -junction-path /sql_hyperv_data -space-guarantee none -percent-snapshot-
space 0

volume create -vserver Hyperv-Work-SVM -volume sql_hyperv_log -aggregate aggr1_node02 -size 1.5TB -
state online -policy default -junction-path /sql_hyperv_log -space-guarantee none -percent-snapshot-
space 0

snapmirror update-ls-set -source-path Hyperv-Work-SVM:rootvol
```

## Create Hyper-V Boot LUNs

To create three boot LUNs for Hyper-V hosts, run the following commands:

```
lun create -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host1 -size 500GB -ostype
windows_2008 -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host2 -size 500GB -ostype
windows_2008 -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host3 -size 500GB -ostype
windows_2008 -space-reserve disabled
```

## Create Storage Volume for SQL Virtual Machines Running on Hyper-V

To create a SQL Virtual Machine Datastore LUN and Swap LUN, run the following commands:

```
lun create -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun sql-vms-csv -size 500GB -ostype
windows_2008 -space-reserve disabled
```

## Create SQL Data and Log LUNs

To create a SQL Virtual Machine Datastore LUN and Swap LUN, run the following commands:

```
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_data -lun SQL-Hyperv-Data-01 -size 400GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_data -lun SQL-Hyperv-Data-02 -size 400GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_data -lun SQL-Hyperv-Data-03 -size 400GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_data -lun SQL-Hyperv-Data-04 -size 400GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_log -lun SQL-Hyperv-Log-01 -size 150GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_log -lun SQL-Hyperv-Log-02 -size 150GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_log -lun SQL-Hyperv-Log-03 -size 150GB -ostype linux -space-reserve disabled
lun create -vserver Hyperv-Work-SVM -volume sql_hyperv_log -lun SQL-Hyperv-Log-04 -size 150GB -ostype linux -space-reserve disabled
```

## Schedule Deduplication

On NetApp AFF systems, deduplication is enabled by default. To schedule deduplication, follow this step:

1. After the volumes are created, assign a once-a-day deduplication schedule to `hyperv_boot` and `infra_datastore_1`:

```
efficiency modify -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -schedule sun-sat@0
efficiency modify -vserver Hyperv-Work-SVM -volume sql_hyperv_data -schedule sun-sat@0
efficiency modify -vserver Hyperv-Work-SVM -volume sql_hyperv_log -schedule sun-sat@0
```

## Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), run the following commands:

```
network interface create -vserver Hyperv-Work-SVM -lif hyperv-iscsi-lif01a -role data -data-protocol iscsi -home-node <st-node01> -home-port a0a-<SQL-HV-iSCSI-A-id> -address <var_node01_hyperv-iscsi-lif01a_ip> -netmask <var_node01_hyperv-iscsi-lif01a_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Hyperv-Work-SVM -lif hyperv-iscsi_lif01b -role data -data-protocol iscsi -home-node <st-node01> -home-port a0a-<SQL-HV-iSCSI-B-id> -address <var_node01_hyperv-iscsi-lif01b_ip> -netmask <var_node01_hyperv-iscsi-lif01b_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Hyperv-Work-SVM -lif hyperv-iscsi_lif02a -role data -data-protocol iscsi -home-node <st-node02> -home-port a0a-<SQL-HV-iSCSI-A-id> -address <var_node02_hyperv-iscsi-lif02a_ip> -netmask <var_node02_hyperv-iscsi-lif02a_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Hyperv-Work-SVM -lif hyperv-iscsi_lif02b -role data -data-protocol iscsi -home-node <st-node02> -home-port a0a-<SQL-HV-iSCSI-B-id> -address <var_node02_hyperv-iscsi-lif02b_ip> -netmask <var_node02_hyperv-iscsi-lif02b_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface show
```

## Add SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, follow these steps:

1. Run the following commands:

```
network interface create -vserver Hyperv-Work-SVM -lif Hyperv-Work-SVM-mgmt -role data -data-protocol none -home-node <st-node01> -home-port e0c -address <Hyperv-Work-SVM-Mgmt_ip> -netmask <Hyperv-Work-SVM-Mgmt_mask> -status-admin up -failover-policy system-defined -firewall-policy mgmt -auto-revert true
```



The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Hyperv-Work-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Hyperv-Work-SVM
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver Hyperv-Work-SVM
```



A cluster serves data through at least one and possibly several SVMs. We have just gone through creating a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create additional SVMs.

## Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into following tables.

**Table 10** iSCSI IQN and iSCSI LIFs

IQN or LIF IP	Info
Hyperv-Work-SVM IQN	
hyperv-iscsi_lif01a_ip	
hyperv-iscsi_lif02a_ip	
hyperv-iscsi_lif01b_ip	
hyperv-iscsi_lif02b_ip	



To obtain the iSCSI IQN, run `iscsi show` command on the storage cluster management interface.

**Table 11** vNIC iSCSI IQNs for Fabric A and Fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
hyperv-host1		<hyperv-host1-iqn>
hyperv-host2		<hyperv-host2-iqn>

Cisco UCS Service Profile Name	iSCSI IQN	Variables
hyperv-host3		<hyperv-host3-iqn>



To obtain the iSCSI vNIC IQN information in the Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then select the iSCSI vNICs tab. The initiator name is displayed at the top of the page under the service profile initiator name.

# Server Configuration

---

## Cisco UCS Base Configuration

This FlexPod deployment will show configuration steps for the Cisco UCS Fabric Interconnects (FI) in a design that will support iSCSI to the NetApp AFF through the Cisco Nexus.

### Perform Initial Setup of Cisco UCS Fabric Interconnects for FlexPod Environments

This section provides detailed procedures for configuring the Cisco Unified Computing System for use in a FlexPod environment. The steps are necessary to provision the Cisco UCS B-Series and C-Series servers and should be followed precisely to avoid improper configuration.

#### Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlexPod environment, follow these steps:

1. Connect to the console port on the first Cisco UCS fabric interconnect.

```
Enter the configuration method: gui
```

```
Physical switch Mgmt0 IP address: <ucsa-mgmt-ip>
```

```
Physical switch Mgmt0 IPv4 netmask: <ucsa-mgmt-mask>
```

```
IPv4 address of the default gateway: <ucsa-mgmt-gateway>
```

2. Using a supported web browser, connect to <https://<ucsa-mgmt-ip>>, accept the security prompts, and click the 'Express Setup' link under HTML.
3. Select Initial Setup and click Submit.
4. Select Enable clustering, Fabric A, and IPv4.
5. Fill in the Virtual IP Address with the UCS cluster IP.
6. Completely fill in the System setup section. For system name, use the overall UCS system name. For the Mgmt IP Address, use <ucsa-mgmt-ip>.
7. Click Submit.

#### Cisco UCS Fabric Interconnect B

To configure the Cisco UCS for use in a FlexPod environment, follow these steps:

1. Connect to the console port on the second Cisco UCS fabric interconnect.

```
Enter the configuration method: gui
```

```
Physical switch Mgmt0 IP address: <ucsb-mgmt-ip>
```

```
Physical switch Mgmt0 IPv4 netmask: <ucsb-mgmt-mask>
```

```
IPv4 address of the default gateway: <ucsb-mgmt-gateway>
```

2. Using a supported web browser, connect to <https://<ucsb-mgmt-ip>>, accept the security prompts, and click the Express Setup link under HTML.

3. Under System setup, enter the Admin Password entered above and click Submit.
4. Enter <ucsb-mgmt-ip> for the Mgmt IP Address and click Submit.

## Cisco UCS Setup

### Log in to Cisco UCS Manager

To log in to the Cisco Unified Computing System environment, follow these steps:

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.



**You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.**

---

2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

### Upgrade Cisco UCS Manager Software to Version 4.0(1c)


This document assumes the use of Cisco UCS 4.0(1c). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(1c), refer to Cisco UCS Manager Install and Upgrade Guides: <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html>

### Anonymous Reporting

To create anonymous reporting, follow this step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products. If you select Yes, enter the IP address of your SMTP Server. Click OK.

Figure 15 Anonymous Reporting



## Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, follow these steps:


1. In Cisco UCS Manager, click Admin.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

## Add Block of IP Addresses for KVM Access

To create a block of IP addresses for in band server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information.

Figure 16 IPv4 IP Addresses for KVM Access



The screenshot shows a dialog box titled "Create Block of IPv4 Addresses". It has a question mark icon and a close button (X) in the top right corner. The dialog contains the following fields:

From :	<input type="text" value="192.168.166.230"/>	Size :	<input type="text" value="20"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="192.168.166.1"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

At the bottom right, there are two buttons: "OK" (blue) and "Cancel" (white with blue border).


5. Click OK to create the block.
6. Click OK in the confirmation message

## Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand All > Time Zone Management.
3. Select Timezone.
4. In the Properties pane, select the appropriate time zone in the Timezone menu.
5. Click Save Changes, and then click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> and click OK. Click OK.
8. Click Add NTP Server.
9. Enter <switch-b-ntp-ip> and click OK. Click OK on the confirmation.

**Figure 17 UCSM NTP Configuration**




### Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS C-Series connectivity. To modify the chassis discovery policy, follow these steps:

1. In Cisco UCS Manager, click Equipment and select Equipment in the second list.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.
5. Click Save Changes.
6. Click OK.

**Figure 18 Edit Discovery Policy**



### Enable Server and Uplink Ports

To enable server and uplink ports, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select Configure as Server Port
5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
7. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.



**The last 6 ports of the Cisco UCS 6332 and Cisco UCS 6332-16UP FIs will only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only**

---

8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

## Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external 2232 FEX modules, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. Expand Chassis and select each chassis that is listed.
3. Right-click each chassis and select Acknowledge Chassis.
4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus 2232 FEX are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and select Acknowledge FEX.
7. Click Yes and then click OK to complete acknowledging the FEX.

## Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.



**In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches**

---

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 153 as the unique ID of the port channel.
6. Enter vPC-153-Nexus as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 154 as the unique ID of the port channel.
16. Enter vPC-154-Nexus as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.



**In the following section, creation of UCSM pools, policies, templates and service profiles will be shown. Create one Sub-Organizations if you planning to have dedicated Organization for your environments. And then create these pools, policies, templates and service profiles under your organization. Else create them in default root organization.**

---

## Create Organization (Optional)

To create an Organization in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profiles > root > Sub-Organizations.
3. Right-click Sub-Organizations and select Create Organization.
4. Enter the name of the Organization and description and click OK.



**For this testing and validation, two organizations are created. One for VMware environment and other for Hyper-V environment.**

---



**While creating the pools, polices, templates and service profiles, navigate to the appropriate organization and create them.**

---



**The pools, polices and templates that are applicable to all the organizations are created under default root organization.**

---

## Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Select Pools > root



**In this procedure, two MAC address pools are created, one for each switching fabric.**

---

3. Right-click MAC Pools under the root.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select Sequential as the option for Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0A:00 as our first MAC address.

---

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server re-sources.
12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click MAC Pools under the root organization.
16. Select Create MAC Pool to create the MAC address pool.
17. Enter MAC-Pool-B as the name of the MAC pool.
18. Optional: Enter a description for the MAC pool.
19. Select Sequential as the option for Assignment Order.
20. Click Next.
21. Click Add.
22. Specify a starting MAC address.



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0B:00 as our first MAC address.

---

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server re-sources.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Figure 19 MAC Pool Address

ID	Assigned
MAC 00:25:B5:B9:0A:00	Yes
MAC 00:25:B5:B9:0A:01	Yes
MAC 00:25:B5:B9:0A:02	Yes
MAC 00:25:B5:B9:0A:03	Yes
MAC 00:25:B5:B9:0A:04	Yes
MAC 00:25:B5:B9:0A:05	Yes
MAC 00:25:B5:B9:0A:06	Yes
MAC 00:25:B5:B9:0A:07	Yes
MAC 00:25:B5:B9:0A:08	Yes
MAC 00:25:B5:B9:0A:09	Yes

ID	Assigned
MAC 00:25:B5:B9:0B:00	Yes
MAC 00:25:B5:B9:0B:01	Yes
MAC 00:25:B5:B9:0B:02	Yes
MAC 00:25:B5:B9:0B:03	Yes
MAC 00:25:B5:B9:0B:04	Yes
MAC 00:25:B5:B9:0B:05	Yes
MAC 00:25:B5:B9:0B:06	Yes
MAC 00:25:B5:B9:0B:07	Yes
MAC 00:25:B5:B9:0B:08	Yes
MAC 00:25:B5:B9:0B:09	Yes

## Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Select Pools > root.
3. Right click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.201-11.com.flexpod as the prefix.
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. Enter ucs-host as the suffix.



**If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.**

12. Enter 1 in the From field.
13. Specify the size of the IQN block sufficient to support the available server resources.
14. Click OK.

15. Click Finish.

## Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Select Pools > root.
3. Right-click IP Pools.
4. Select Create IP Pool.
5. Enter SQL-FP-iSCSI-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP address.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses
11. Set the size to enough addresses to accommodate the servers
12. Click OK.
13. Click Next.
14. Click Finish.
15. Right-click IP Pools.
16. Select Create IP Pool.
17. Enter SQL-FP-iSCSI-Pool-A as the name of IP pool.
18. Optional: Enter a description for the IP pool.
19. Select Sequential for the assignment order.
20. Click Next.
21. Click Add to add a block of IP address.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers
24. Click OK.
25. Click Next.

26. Click Finish.

## Create UUID Suffix Pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

## Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, follow these steps:



**Consider creating unique server pools to achieve the granularity that is required in your environment. For instance, create separate server pools for VMware and Hyper-V Cluster in their corresponding organizations.**

---

1. In Cisco UCS Manager, click Servers.
2. Select Pools > root > Sub-Organizations >
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter SQL-FP-Blades-Pool as the name of the server pool.

6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware/Hyper-V cluster and click >> to add them to the SQL-FP-Blades-Pool server pool.
9. Click Finish.
10. Click OK.

## Create VLANs

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.



**In this procedure, ten unique VLANs are created, Four for VMware cluster and six for Hyper-V cluster. See Table 2**

---

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.
10. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs
14. Enter SQL-VM-MGMT as the name of the VLAN to be used for VMware management traffic.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the VMware In-Band management VLAN ID.
17. Keep the Sharing Type as None.

18. Click OK, and then click OK again.
19. Right-click VLANs.
20. Select Create VLANs.
21. Enter SQL-VM-iSCSI-A as the name of the VLAN to be used for VMware iSCSI traffic.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the VMWare iSCSI-A VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter SQL-VM-iSCSI-B as the name of the VLAN to be used for VMware iSCSI traffic.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the VMWare iSCSI-B VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK, and then click OK again.
33. Right-click VLANs.
34. Select Create VLANs.
35. Enter SQL-vMotion as the name of the VLAN to be used for vMotion.
36. Keep the Common/Global option selected for the scope of the VLAN.
37. Enter the SQL-vMotion VLAN ID.
38. Keep the Sharing Type as None.
39. Click OK, and then click OK again.
40. Select Create VLANs.
41. Enter SQL-HV-MGMT as the name of the VLAN to be used for Hyper-V Management traffic.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the Hyper-V In-Band Mgmt VLAN ID.
44. Keep the Sharing Type as None.

45. Click OK, and then click OK again.
46. Right-click VLANs.
47. Select Create VLANs.
48. Enter SQL-HV-iSCSI-A as the name of the VLAN to be used for Hyper-V iSCSI traffic.
49. Keep the Common/Global option selected for the scope of the VLAN.
50. Enter the Hyper-V iSCSI-A VLAN ID.
51. Keep the Sharing Type as None.
52. Click OK, and then click OK again.
53. Right-click VLANs.
54. Select Create VLANs.
55. Enter SQL-HV-iSCSI-B as the name of the VLAN to be used for Hyper-V iSCSI traffic.
56. Keep the Common/Global option selected for the scope of the VLAN.
57. Enter the Hyper-V iSCSI-B VLAN ID.
58. Keep the Sharing Type as None.
59. Click OK and then click OK again.
60. Right-click VLANs.
61. Select Create VLANs.
62. Enter SQL-Live-Migration as the name of the VLAN to be used for Hyper-V live migration traffic.
63. Keep the Common/Global option selected for the scope of the VLAN.
64. Enter the Hyper-V live migration VLAN ID.
65. Keep the Sharing Type as None.
66. Click OK, and then click OK again.
67. Right-click VLANs.
68. Select Create VLANs.
69. Enter SQL-HV-Cluster as the name of the VLAN to be used for Hyper-V cluster traffic.
70. Keep the Common/Global option selected for the scope of the VLAN.
71. Enter the Hyper-V SQL-HV-Cluster VLAN ID.

72. Keep the Sharing Type as None.

73. Click OK and then click OK again.


### Modify Default Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root > <Sub-Organization>
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 4.0(1c) for both the Blade and Rack Packages.
7. Click OK then click OK again to modify the host firmware package.

Figure 20 Host Firmware Policy



## Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes.
6. Click OK.

**Figure 21 QoS Policy**

LAN / LAN Cloud / QoS System Class

General Events FSM

Actions Properties

Use Global Owner: Local

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

## Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.



**This policy should not be used on servers that contain local disks.**

To create a local disk configuration policy, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter SAN-Boot as the local disk configuration policy name.
6. Change the mode to No Local Storage.

7. Click OK to create the local disk configuration policy.
8. Click OK.

**Figure 22 Local Disk Policy**

**Create Local Disk Configuration Policy** [?] [X]

Name : No-Local-Disk

Description : Nothing is on the local disk

Mode : No Local Storage

**FlexFlash**

FlexFlash State :  Disable  Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.  
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State :  Disable  Enable

FlexFlash Removable State :  Yes  No  No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily.  
Please ensure SD cards are not in use before changing the FlexFlash Removable State.

OK Cancel

## Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)

To create a network control policy that enables CDP and LLDP on virtual network ports, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

9. Click OK.

Figure 23 Network Control Policy

**Create Network Control Policy** ? X

CDP :  Disabled  Enabled

MAC Register Mode :  Only Native Vlan  All Host Vlans

Action on Uplink Fail :  Link Down  Warning

**MAC Security**

Forge :  Allow  Deny

**LLDP**

Transmit :  Disabled  Enabled

Receive :  Disabled  Enabled

OK Cancel

## Create Power Control Policy

To create a power control policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Figure 24 Power Policy

**Create Power Control Policy** ? X

Name : SQL-VMNoPowerCAP

Description : No Power CAP on ESXi blades hosting SQLVMs

Fan Speed Policy : Any

**Power Capping**

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap  cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK Cancel

## Create Server BIOS Policy


To create a server BIOS policy for the Cisco UCS environment, follow these steps:



**Make sure to carefully consider the bios settings you want to use for your workloads. For latency intensive and critical database workloads, the following bios settings are recommended. BIOS settings may also vary based on the virtualization platforms (VMware ESXi/Hyper-V).**

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter Virtual-Host as the BIOS policy name and enter description (optional)
6. Click OK.
7. Expand the BIOS Policies and select Virtual-Host policy.
8. In the right pane, Change CDN Control setting to Enabled.
9. Change the Quiet Boot setting to disabled.

Figure 25 BIOS Policy-1



10. Click the Advanced tab > Processor tab and click BIOS Settings to list the settings in ascending order.

11. Change the following settings as follows:

- a. Adjacent Cache Line Prefetcher -> Enabled
- b. Boot Performance Mode -> Max Performance
- c. CPU Hardware Power Management -> HWPM Native Mode
- d. CPU Performance -> Enterprise
- e. DCU IP Prefetcher -> Enabled
- f. DCU Streamer Prefetch -> Enabled
- g. DRAM Clock Throttling -> Performance
- h. Energy Performance -> Performance
- i. Energy Performance Tuning -> OS
- j. Enhance Intel Speed Step Tech -> Enabled
- k. Frequency Floor Override -> Enabled
- l. Hardware Prefetcher -> Enabled
- m. IMC interleave -> 1-way Interleave
- n. Intel Hyper Threading Tech -> Enabled
- o. Intel Virtualization Tech -> Enabled
- p. LLC Prefetch -> Disabled
- q. P State Coordination -> HW ALL
- r. Package C State Limit -> C0 C1 State
- s. Patrol Scrub -> Disabled

- t. Power Technology -> Performance
- u. Processor C State -> Disabled
- v. Processor C1E -> Disabled
- w. Processor C3 -> Disabled
- x. Processor C6 -> Disabled
- y. Processor C7 -> Disabled
- z. Sub NUMA Clustering -> Enabled
- aa. UPI Prefetch -> Enabled
- bb. XPT Prefetch -> Disabled

Figure 26 BIOS Policy-2

The screenshot shows a web-based BIOS configuration interface. The breadcrumb trail is: Servers / Policies / root / Sub-Organizations / SQL-FP / BIOS Policies / Virtual-Host. The 'Advanced' tab is selected, and the 'Processor' sub-tab is active. The interface displays a table of BIOS settings with columns for the setting name and its current value. At the bottom, there are 'Save Changes' and 'Reset Values' buttons. A watermark 'Activate Windows' is visible in the bottom right corner.

BIOS Setting	Value
AMD Memory Interleaving	Platform Default
AMD Memory Interleaving Size	Platform Default
Adjacent Cache Line Prefetcher	Enabled
Altitude	Platform Default
Autonomous Core C-state	Platform Default
Bank Group Swap	Platform Default
Boot Performance Mode	Max Performance
CPU Hardware Power Management	HWPM Native Mode
CPU Performance	Enterprise
Channel Interleaving	Platform Default
Chipselect Interleaving	Platform Default
Core Multi Processing	Platform Default

Figure 27 BIOS Policy-3

Servers / Policies / root / Sub-Organizations / SQL-VMWare-Clus / BIOS Policies / SQL-VMWare-Host

Main **Advanced** Boot Options Server Management Events

**Processor** Intel Directed IO RAS Memory Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
Core Multi Processing	Platform Default
Core Performance Boost	Platform Default
DCU IP Prefetcher	Enabled
DCU Streamer Prefetch	Enabled
DRAM Clock Throttling	Performance
Demand Scrub	Platform Default
Determinism Slider	Platform Default
Direct Cache Access	Platform Default
Downcore control	Platform Default
Energy Efficient Turbo	Platform Default
Energy Performance	Performance
Energy Performance Tuning	OS

Figure 28 BIOS Policy-4

Servers / Policies / root / Sub-Organizations / SQL-VMWare-Clus / BIOS Policies / SQL-VMWare-Host

Main **Advanced** Boot Options Server Management Events

**Processor** Intel Directed IO RAS Memory Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
Energy Performance Tuning	OS
Enhanced Intel SpeedStep Tech	Enabled
Execute Disable Bit	Platform Default
Frequency Floor Override	Enabled
Global C-state Control	Platform Default
Hardware Prefetcher	Enabled
IMC Inteleave	1-way Interleave
IOMMU	Platform Default
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	Platform Default
Intel Virtualization Technology	Enabled
L1 Stream HW Prefetcher	Platform Default

Figure 29 BIOS Policy-5

Servers / Policies / root / Sub-Organizations / SQL-VMWare-Clus / BIOS Policies / SQL-VMWare-Host

Main | **Advanced** | Boot Options | Server Management | Events

**Processor** | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform | Graphics Configuration

Advanced Filter | Export | Print

BIOS Setting	Value
L1 Stream HW Prefetcher	Platform Default
L2 Stream HW Prefetcher	Platform Default
LLC Prefetch	Disabled
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
Memory Interleaving	Platform Default
P-STATE Coordination	HW ALL
Package C-State Limit	C0 C1 State
Patrol Scrub	Disabled
Power Technology	Performance
Processor C-State	Disabled
Processor C1E	Disabled

Figure 30 BIOS Policy-6

Servers / Policies / root / Sub-Organizations / SQL-VMWare-Clus / BIOS Policies / SQL-VMWare-Host

Main | **Advanced** | Boot Options | Server Management | Events

**Processor** | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform | Graphics Configuration

Advanced Filter | Export | Print

BIOS Setting	Value
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMCi	Platform Default
ProcessorEppProfile	Platform Default
Rank Interleaving	Platform Default
SMEE	Platform Default
SMT Mode	Platform Default
SVM Mode	Platform Default
Sub NUMA Clustering	Enabled
UPI Prefetch	Enabled

Workload Configuration	Platform Default
XPT Prefetch	Disabled

Add | Delete | Info


12. Click the Advanced tab > RAS Memory tab and Change LV DDR mode to Performance Mode.

13. Click Save Changes.

### Create Maintenance Policy

To update new Maintenance Policy, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root.
3. Select Maintenance Policies > right-click > Select Create Maintenance policy.
4. Enter name of the Maintenance Policy.
5. Change Storage Config Deployment Policy to User Ack.
6. Change the Reboot Policy to User Ack.
7. Select "On Next Boot" to delegate maintenance windows to server administrators.
8. Click Save Changes.
9. Click OK to accept the change.

**Figure 31 Maintenance Mode**

**Create Maintenance Policy**

Name : Hots-Maintenance

Description :

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy :  Immediate  User Ack

Reboot Policy :  Immediate  User Ack  Timer Automatic

On Next Boot (Apply pending changes at next reboot.)

OK Cancel

## Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, follow these steps. A total of 4 vNIC Templates will be created.

### Create Infrastructure vNICs

1. In Cisco UCS Manager, click LAN.
2. Select Policies > root > Sub-Organization.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.

5. Enter SQL-FP-UPLINK-A as the vNIC template name and enter description.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the appropriate VLANs for your environment
  - a. For VMware environment, select check boxes for SQL-VM-MGMT, SQL-vMotion, SQL-VM-iSCSI-A, SQL-VM-iSCSI-B and Native-VLAN.
  - b. For Hyper-V Cluster, select check boxes for SQL-HV-MGMT, SQL-Live-Migration, SQL-HV-iSCSI-A, SQL-HV-iSCSI-B, SQL-HV-Cluster and Native-VLAN
13. Set Native-VLAN as the native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-Pool-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.
18. Click OK to create the vNIC template.
19. Click OK.

Figure 32 vNIC Template for Host Infrastructure Traffic (VMware)

**Create vNIC Template**

Name : SQL-FP-UPLINK-A  
 Description : UPLINK 1  
 Fabric ID :  Fabric A  Fabric B  Enable Failover

Redundancy

Redundancy Type :  No Redundancy  Primary Template  Secondary Template  
 Peer Redundancy Template : <not set>

**Target**

Adapter  
 VM

Warning

If **VM** is selected, a port profile by the same name will be created.  
 If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type :  Initial Template  Updating Template

**VLANs** | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	Infra-vMotion	<input type="radio"/>
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	SQL-VM-ISCSI-A	<input type="radio"/>
<input checked="" type="checkbox"/>	SQL-VM-ISCSI-B	<input type="radio"/>
<input checked="" type="checkbox"/>	SQL-VM-MGMT	<input type="radio"/>
<input checked="" type="checkbox"/>	SQL-vMotion	<input type="radio"/>

Create VLAN

CDN Source :  vNIC Name  User Defined  
 MTU : 9000  
 MAC Pool : MAC-Pool-A(196/200)  
 QoS Policy : <not set>  
 Network Control Policy : Enable-CDP-LLDP  
 Pin Group : <not set>

Connection Policies

Dynamic vNIC  usNIC  VMQ  
 Dynamic vNIC Connection Policy : <not set>

Create the Secondary Redundancy Template SQL-FP-UPLINK-B

To create the secondary redundancy template SQL-FP-UPLINK-B, follow these steps:

1. Select LAN.
2. Select Policies > root > Sub-Organization.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.

5. Enter SQL-FP-UPLINK-B as the vNIC template name.
6. Select Fabric B.
7. Do not select the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select SQL-FP-UPLINK-A for the Peer Redundancy Template.
10. In the MAC Pool list, select MAC-Pool-B. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Select the rest of the options as you entered for SQL-FP-UPLINK-A .
12. Click OK to create the vNIC template.
13. Click OK.

### Create iSCSI vNICs

To create iSCSI vNICs, follow these steps:

1. Select LAN.
2. Select Policies > root > <Sub Organization> (Optional).
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter iSCSI-Boot-A as the vNIC template name and enter description.
6. Select Fabric A. Do not select the Enable Failover checkbox.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter checkbox is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only SQL-VM-iSCSI-A for VMware Cluster. SQL-HV-iSCSI-A for Hyper-V Cluster.
11. Select iSCSI-A-VLAN/ SQL-HV-iSCSI-A as the native VLAN.
12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select MAC-Pool-A.
15. From the Network Control Policy list, select Enable-CDP-LLDP.
16. Click OK to complete creating the vNIC template

17. Click OK.
18. Select LAN.
19. Select Policies > root > <Sub-Organization> (Optional).
20. Right-click vNIC Templates.
21. Select Create vNIC Template.
22. Enter iSCSI-Template-B as the vNIC template name.
23. Select Fabric B. Do not select the Enable Failover checkbox.
24. Leave Redundancy Type set at No Redundancy.
25. Under Target, make sure that only the Adapter checkbox is selected.
26. Select Updating Template for Template Type.
27. Under VLANs, select only SQL-VM-iSCSI-B for VMware Cluster. SQL-HV-iSCSI-B for Hyper-V Cluster.
28. Select SQL-VM-iSCSI-B/ SQL-HV-iSCSI-B as the native VLAN.
29. Leave vNIC Name set for the CDN Source.
30. Under MTU, enter 9000.
31. From the MAC Pool list, select MAC-Pool-B.
32. From the Network Control Policy list, select Enable-CDP-LLDP.
33. Click OK to complete creating the vNIC template.
34. Click OK.

### Creating Adapter Policy

The default 'vmware' adapter policy is not configured with sufficient transmit and receive queues. For high IO demanding SQL database deployments, it is recommended to customized adapter policy. To create adapter policy VMware environments, follow these steps:



**For Windows environments, use the default 'windows' adapter policy as it is configured enough transmit and receive queues.**

---

1. In Cisco UCS Manager, click LAN.
2. Select Servers > Policies > root > <Sub-Organization> (optional).
3. Right-click Adapter Policy and select Create Ethernet Adapter Policy.
4. Enter name as VMware-HighTrf and enter Description.
5. Enter the settings as shown in the following screenshot.

Figure 33 Adapter Policy

Servers / Policies / root / Sub-Organizations / SQL-VMWare-Clus / Adapter Policies / Eth Adapter Policy...

General Events

---

**Actions**

Delete

Show Policy Usage

Use Global

**Properties**

Name : **VMware-HighTrf**

Description : Adapter policy for ESXi

Owner : **Local**

---

**Resources**

Pooled :  Disabled  Enabled

Transmit Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

---

Receive Queues : 8 [1-1000]

Ring Size : 4096 [64-4096]

---

Completion Queues: 16 [1-2000]

Interrupts : 18 [1-1024]

---

**Options**

Transmit Checksum Offload :  Disabled  Enabled

Receive Checksum Offload :  Disabled  Enabled

TCP Segmentation Offload :  Disabled  Enabled

TCP Large Receive Offload :  Disabled  Enabled

Receive Side Scaling (RSS) :  Disabled  Enabled

Accelerated Receive Flow Steering :  Disabled  Enabled

Network Virtualization using Generic Routing Encapsulation :  Disabled  Enabled

Virtual Extensible LAN :  Disabled  Enabled

Failback Timeout (Seconds) : 5 [0-600]

Interrupt Mode :  MSI X  MSI  IN Tx

Interrupt Coalescing Type :  Min  Idle

Interrupt Timer (us) : 125 [0-65535]

RoCE :  Disabled  Enabled

Advance Filter :  Disabled  Enabled

Interrupt Scaling :  Disabled  Enabled

### Create LAN Connectivity Policy for iSCSI Boot

To configure the necessary Infrastructure LAN Connectivity Policy, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Select LAN > Policies > root > Sub-Organization.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter iSCSI-Boot as the name of the policy.

6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter 00-Host-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select SQL-FP-UPLINK-A.
10. In the Adapter Policy list, select VMWare-HighTrf for VMware environment and Windows for Hyper-V environment.
11. Click OK to add this vNIC to the policy.

**Figure 34 LAN Connectivity Policy-1**

The screenshot shows a 'Create vNIC' dialog box with the following configuration:

- Name: 00-Host-A
- Use vNIC Template:
- Redundancy Pair:
- vNIC Template: SQL-FP-UPLINK-A
- Adapter Policy: VMware-HighTrf
- Peer Name: (empty)
- Buttons: Create vNIC Template, Create Ethernet Adapter Policy

12. Click Add to add another vNIC to the policy.
13. In the Create vNIC box, enter 01-Host-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select SQL-FP-UPLINK-B.
16. In the Adapter Policy list, select VMWare-HighTrf for VMware environment and Windows for Hyper-V environment.
17. Click OK to add the vNIC to the policy.
18. Click the upper Add button to add a vNIC.
19. In the Create vNIC dialog box, enter 02-iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI-Boot-A.
22. In the Adapter Policy list, select VMWare for VMware environment and Windows for Hyper-V environment.
23. Click OK to add this vNIC to the policy.
24. Click the upper Add button to add a vNIC to the policy.

25. In the Create vNIC dialog box, enter 03-iSCSI-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI-Boot-B.
28. In the Adapter Policy list, select VMWare for VMware environment and Windows for Hyper-V environment.
29. Click OK to add this vNIC to the policy.
30. Expand the Add iSCSI vNICs.
31. Select Add in the Add iSCSI vNICs section.
32. Set the name to iSCSI-A-vNIC.
33. Select the 02-iSCSI-A as Overlay vNIC.
34. Depending on the environment, Set the VLAN to SQL-VM-iSCSI-A / SQL-HV-iSCSI-A (native).
35. Set the iSCSI Adapter Policy to default.
36. Leave the MAC Address set to None.
37. Click OK.
38. Select Add in the Add iSCSI vNICs section.
39. Set the name to iSCSI-B-vNIC.
40. Select the 03-iSCSI-A as Overlay vNIC.
41. Depending on the environment, Set the VLAN to SQL-VM-iSCSI-A / SQL-HV-iSCSI-A (native).
42. Set the iSCSI Adapter Policy to default.
43. Leave the MAC Address set to None.
44. Click OK, then click OK again to create the LAN Connectivity Policy.

Figure 35 LAN Connectivity Policy-2

Create LAN Connectivity Policy
? X

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Native VLAN
vNIC 03-ISCSI-B	Derived	
vNIC 02-ISCSI-A	Derived	
vNIC 01-Host-B	Derived	
vNIC 00-Host-A	Derived	

Delete + Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
ISCSI vNIC ISCSI-B-vNIC	03-ISCSI-B		Derived
ISCSI vNIC ISCSI-A-vNIC	02-ISCSI-A		Derived

+ Add Delete Modify

OK Cancel

## Create vMedia Policy for installing both ESXi and Hyper-V Hypervisors

For this CVD, a Linux-based host is used to provide the HTTP webserver with the required ISO images for vMedia policies to connect and to install hypervisors.



**Make sure to download the required ISOs and upload them to the webserver for the HTTP file transfer.**

The vMedia Policy created will map the hypervisor images (ISOs) to the Cisco UCS servers to boot from and install. To create this policy for both the hypervisor environments, follow these steps:

1. In Cisco UCS Manager, select Servers.
2. Select Policies > root > <Sub-Organization> (Optional).
3. Right-click vMedia Policies.
4. Select Create vMedia Policy.
5. Provide a name for the policy.
6. Enter Description as appropriate description.
7. Click Add.
8. Enter a name for mount.

9. Select the CDD Device Type.
10. Select the HTTP Protocol.
11. Enter the IP Address of the web server.

---

 **Since the DNS server IPs were not entered into the KVM IP earlier, it is necessary to enter the IP of the web server instead of the hostname.**

---

12. Enter the name of the ISO files that you are using for installation as the Remote File name. For example, enter VMware\_ESXi\_6.7.0\_10302608\_Custom\_Cisco\_6.7.1.1.iso for VMware environments.
13. Enter the web server path to the ISO file in the Remote Path field.
14. Click OK to create the vMedia Mount.
15. Click OK then OK again to complete creating the vMedia Policy.


---

 **For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi/Windows host. On first boot the host will boot into the ESXi/Windows installer since the SAN mounted disk is empty. After ESXi/Windows is installed, the vMedia will not be referenced as long as the boot disk is accessible.**

---

Figure 36 shows the sample vMedia mount for VMware environment.

**Figure 36 vMedia Policy**



**Create vMedia Mount** [?] [X]

Name : ESX6.7-HTTP

Description :

Device Type :  CDD  HDD

Protocol :  NFS  CIFS  HTTP  HTTPS

Hostname/IP Address : 192.168.166.150

Image Name Variable :  None  Service Profile Name

Remote File : VMware\_ESXi\_6.7.0\_10302608\_Custom\_Cisco\_6

Remote Path : /software/vSphere6.7/

Username :

Password :

Remap on Eject :

OK Cancel


## Create Boot Policy for iSCSI Boot

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi\_lif01a and iscsi\_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi\_lif02a and iscsi\_lif02b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).

To create a boot policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Policies > root > <Sub-Organization> (Optional).
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter VMHost-iSCSIBoot as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Select Uefi for Boot Mode.
9. Select the Boot Security Check Box.
10. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
11. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
12. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC and Click OK.
13. Select iSCSI-A-vNIC and click on Set Uefi Boot Parameters.
14. Set Boot Loader Name to 'BOOTx64.EFI'
15. Set Boot Loader path: \EFI\BOOT\
16. Click OK.

Figure 37 Uefi Boot Configuration



Set Uefi Boot Parameters

Uefi Boot Parameters

Boot Loader Name : BOOTx64.EFI

Boot Loader Path : \\EFI\\BOOT\\

Boot Loader Description :

OK Cancel



For VMware environment set Boot Loader Path \\EFI\\BOOT\\. For Windows environment set Boot Loader Path to '/efi/boot/'

17. Add iSCSI Boot.
18. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
19. Click OK.
20. Select iSCSI-B-vNIC and click on Set Uefi Boot Parameters.
21. Set Boot Loader Name to 'BOOTx64.EFI'
22. Set Boot Loader path: \\EFI\\BOOT\\
23. Click OK.
24. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.
25. Click OK to create the policy.

Figure 38 Boot Policy

**Create Boot Policy**

Name : VMHost-iSCSIBoot

Description : UEFI Boot policy for Host to boot from iSCSI

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode :  Legacy  Uefi

Boot Security :

**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.  
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

Add iSCSI Boot

**Boot Order**

Name	vNIC/vHBA/iSCSI...	Type	LU...	W...	Slo...	Bo...	Bo...	Des...
Remote CD/DVD	1							
iSCSI								
ISCSI	ISCSI-A-vNIC	Primary						
ISCSI	ISCSI-B-vNIC	Secondary						

Move Up Move Down Delete

OK Cancel

## Create Service Profile Templates

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Select Service Profile Templates > root > <Sub-Organization> (Optional).
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter SQL-FP-Host-iSCSIBoot as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the "Updating Template" option.
7. Under UUID, select UUID\_Pool as the UUID pool.
8. Click Next.

Figure 39 Create Service Profile Template

**Create Service Profile Template** ? X

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.  
Where : **org-root/org-SQL-FP**

The template will be created in the following organization. Its name must be unique within this organization.  
Type :  Initial Template  Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.  
UUID

UUID Assignment:

The UUID will be assigned from the selected pool.  
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev    Next >    **Finish**    Cancel

### Configure Storage Provisioning

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy and select the No-Local-Disk Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

### Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select "Use the Connectivity Policy" option to configure the LAN connectivity.
3. Select iSCSI-Boot from the LAN Connectivity Policy drop-down list.
4. Select IQN\_Pool in Initiator Name Assignment.
5. Click Next.

Figure 40 Service Profile Template: Networking

The screenshot shows the 'Create Service Profile Template' wizard in the Networking step. The left sidebar lists 11 steps: 1. Identify Service Profile Template, 2. Storage Provisioning, 3. Networking (selected), 4. SAN Connectivity, 5. Zoning, 6. vNIC/HBA Placement, 7. vMedia Policy, 8. Server Boot Order, 9. Maintenance Policy, 10. Server Assignment, and 11. Operational Policies.

The main configuration area includes:

- Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) [Dropdown]
- How would you like to configure LAN connectivity?
  - Simple
  - Expert
  - No vNICs
  - Use Connectivity Policy
- LAN Connectivity Policy : ISCSI-Boot [Dropdown]
- Initiator Name Assignment: IQN-Pool(15/20) [Dropdown]
- Initiator Name : [Create IQN Suffix Pool](#)
- The IQN will be assigned from the selected pool. The available/total IQNs are displayed after the pool name.

Navigation buttons at the bottom: < Prev, Next >, Finish, Cancel.

### Configure Storage Options

1. Select No vHBAs for the "How would you like to configure SAN connectivity?" field.
2. Click Next.

### Configure Zoning Options

1. Click Next.

### Configure vNIC/HBA Placement

1. In the "Select Placement" list, leave the placement policy as "Let System Perfmon Placement".
2. Click Next.

### Configure vMedia Policy

1. Do not select a vMedia Policy.
2. Click Next.

### Configure Server Boot Order

1. Select VMHost-iSCSIBoot for Boot Policy.
2. In the Boot order, select iSCSI-A-vNIC.
3. Click Set iSCSI Boot Parameters button.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment.

5. Leave the "initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
6. Set SQL-FP-iSCSI-Pool-A initiator IP Address Policy
7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI Target Name of SVM for your environment. To get the iSCSI target name of the corresponding SVM, login into storage cluster management interface and run "iscsi show" command.

**Figure 41 iSCSI Target IQN**

```
bb09-affa300::>
bb09-affa300::> iscsi show

```

Vserver	Target Name	Target Alias	Status Admin
ESXi-Work-SVM	<u>iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5</u>	bb09-Work-SVM	up
Hyperv-Work-SVM	<u>iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7</u>	Hyperv-Work-SVM	up
bb09-Infra-SVM	iqn.1992-08.com.netapp:sn.7005a076f32b11e8a3ac00a098a9fed2:vs.3	bb09-Infra-SVM	up

```

3 entries were displayed.
bb09-affa300::> █

```

10. Enter the IP address of iscsi\_lif\_02a for the IPv4 Address field.

**Figure 42 Setting iSCSI Target IP**

**Create iSCSI Static Target** ? X

iSCSI Target Name :

Priority :

Port :

Authentication Profile :  [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

11. Click OK to add the iSCSI static target.



For Hyper-V deployments, add only one target IP for iSCSI-A-vNIC now and proceed with the next item; configuring Maintenance Policy. Avoid entering a second target IP. Only after installing the MPIO feature on the Hyper-V hosts, return to this section and add the remaining paths and targets as shown in the following steps. This step is to avoid data corruption before any multipath drivers are in place. For the VMware environment, it is okay to proceed with the following steps.

12. Click Add.
13. Enter the iSCSI Target Name.
14. Enter the IP address of iscsi\_lif\_01a for the IPv4 Address field.
15. Click OK to add the iSCSI static target.

**Figure 43 Set iSCSI Boot Parameter-1**

### Set iSCSI Boot Parameters ? X

**Initiator Address**

Initiator IP Address Policy: SQL-FP-iSCSI-POOL-A(8/10) ▾

IPv4 Address : **0.0.0.0**  
 Subnet Mask : **255.255.255.0**  
 Default Gateway : **0.0.0.0**  
 Primary DNS : **0.0.0.0**  
 Secondary DNS : **0.0.0.0**

[Create IP Pool](#)  
 The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface  iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPv4 Addre...	LUN Id
iqn.1992-08...	1	3260		192.168.13.19	0
iqn.1992-08...	2	3260		192.168.13.18	0

⊕ Add 🗑 Delete ℹ Info

**Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.**

OK
Cancel




The target IPs were added with the storage Node 02 IP first and the storage Node 01 IP second. This is assuming the boot LUN is on Node 01. The host will boot using the path to Node 01 if the order in this procedure is used.

16. In the Boot order, select iSCSI-B-vNIC.

17. Click Set iSCSI Boot Parameters.
18. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment
19. Leave the "initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps
20. Set SQL-FP-iSCSI-POOL-B as the "initiator IP Address Policy"
21. Select iSCSI Static Target Interface option.
22. Click Add.
23. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-SVM, login into storage cluster management interface and run "iscsi show" command.
24. Enter the IP address of iscsi\_lif\_02b for the IPv4 Address field.
25. Click OK to add the iSCSI static target.
26. Click Add.
27. Enter the iSCSI Target Name.
28. Enter the IP address of iscsi\_lif\_01b for the IPv4 Address field.
29. Click OK to add the iSCSI static target.

**Figure 44 Set iSCSI Boot Parameters-2**




30. Click Next.

### Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to Host-Maintenance.

**Figure 45 Host Maintenance Policy**



2. Click Next.

### Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, select SQL-FP-Blades.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Expand Firmware Management at the bottom of the page and select the default policy.

Figure 46 Server Assignment

The screenshot shows the 'Create Service Profile Template' wizard. The left sidebar has 11 steps, with '10 Server Assignment' highlighted. The main content area is titled 'Create Service Profile Template' and includes the following elements:

- A header with a help icon (?) and a close icon (X).
- A text box: 'Optionally specify a server pool for this service profile template.'
- Text: 'You can select a server pool you want to associate with this service profile template.'
- 'Pool Assignment:' dropdown menu set to 'SQL-FP-Blades' and a 'Create Server Pool' link.
- Text: 'Select the power state to be applied when this profile is associated with the server.'
- Radio buttons for 'Up' and 'Down', with 'Down' selected.
- Text: 'The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.'
- 'Server Pool Qualification:' dropdown menu set to '<not set>'
- 'Restrict Migration:' checkbox, which is unchecked.
- A section titled 'Firmware Management (BIOS, Disk Controller, Adapter)' with a minus sign icon.
- Text: 'If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.'
- 'Host Firmware Package:' dropdown menu set to 'default' and a 'Create Host Firmware Package' link.
- Navigation buttons at the bottom: '< Prev', 'Next >', 'Finish' (highlighted in blue), and 'Cancel'.

4. Click Next.

### Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, select Virtual-Host.
2. Expand Power Control Policy Configuration and select Host-No-PowerCap in the Power Control Policy list.
3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

### Create vMedia-Enabled Service Profile Template

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Select Service Profile Templates > root > <Sub-Organization>.
3. Service Template SQL-FP-Host-iSCSIBoot.
4. Right-click SQL-FP-Host-iSCSIBoot and select Create a Clone.
5. Name the clone SQL-FP-Host-iSCSIBoot-vM.
6. Select the newly-created SQL-FP-Host-iSCSIBoot-vM and select the vMedia Policy tab.
7. Click Modify vMedia Policy.
8. Select the appropriate vMedia policy previously created and click OK.


9. Click OK to confirm.

## Create Service Profiles

To create service profiles from the service profile template, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Select Service Profile Templates > root > Sub-Organization (Optional).
3. Service Template SQL-FP-Host-iSCSIBoot-VM.
4. Right-click SQL-FP-Host-iSCSIBoot-VM and select Create Service Profiles from Template.
5. Enter VM-Host-0 as the service profile prefix.
6. Enter 1 as the Name Suffix Starting Number.
7. Enter a number for the required Number of Instances.
8. Click OK to create the service profiles.

**Figure 47** Creating Service Profiles



The screenshot shows a dialog box titled "Create Service Profiles From Template". It has a title bar with a question mark icon and a close button (X). The dialog contains three input fields:

- Naming Prefix : VM-Host-0
- Name Suffix Starting Number : 1
- Number of Instances : 2

At the bottom right, there are two buttons: "OK" (blue) and "Cancel" (white with a grey border).

9. Click OK in the confirmation message.
10. When VMware ESXi 6.7/Windows server 2016 has been installed on the hosts, the host Service Profiles can be bound to the VM-SQL-FP-Host-iSCSIBoot Service Profile Template to remove the vMedia Mapping from the host.

## Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

## Gather Necessary Information

After the Cisco UCS service profiles are created, each hypervisor host in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 12 and Table 13

**Table 12 iSCSI LIFs for iSCSI IQN**

SVM	Target: IQN
ESXi-Work-SVM	
Hyperv-Work-SVM	



To obtain the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

**Table 13 vNIC iSCSI IQNs for Fabric A and Fabric B**

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-01		<vm-host-01-iqn>
VM-Host-02		<vm-host-02-iqn>
VM-Host-03		<vm-host-03-iqn>



To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to **Servers > Service Profiles > root > Sub-Organization**. Click on each service profile and then click the "iSCSI vNICs" tab. The "Initiator Name" is displayed at the top of the page under the "Service profile Initiator Name."

## ONTAP Boot, Data LUNs, and Igroups Setup

### Create igroups

#### VMware Cluster

To create igroups for the VMware cluster, follow these steps:

1. To create igroups, run the following commands from the cluster management node SSH connection.

```
igroup create -vserver ESXi-Work-SVM -igroup esxi-host1 -protocol iscsi -ostype vmware -initiator
<esxi-host1-iqn>
igroup create -vserver ESXi-Work-SVM -igroup esxi-host2 -protocol iscsi -ostype vmware -initiator
<esxi-host2-iqn>
igroup create -vserver ESXi-Work-SVM -igroup esxi-host3 -protocol iscsi -ostype vmware -initiator <
esxi-host3-iqn>

igroup create -vserver ESXi-Work-SVM -igroup esxi-Hosts -protocol iscsi -ostype vmware -initiator
<esxi-host1-iqn>,<esxi-host2-iqn>,<esxi-host3-iqn>
```

2. To view the three new igroups, run the `igroup show` command.

#### Hyper-V Cluster

To create igroups for the Hyper-V cluster, follow these steps:

1. To create igroups, run the following commands from the cluster management node SSH connection.

```
igroup create -vserver Hyperv-Work-SVM -igroup hyperv-host1 -protocol iscsi -ostype Windows_2008 -
initiator <hyperv-host1-iqn>
igroup create -vserver Hyperv-Work-SVM -igroup hyperv-host2 -protocol iscsi -ostype Windows_2008 -
initiator <hyperv-host2-iqn>
igroup create -vserver Hyperv-Work-SVM -igroup hyperv-host3 -protocol iscsi -ostype Windows_2008 -
initiator < hyperv-host3-iqn>

igroup create -vserver Hyperv-Work-SVM -igroup hyperv-hosts -protocol iscsi -ostype vmware -initiator
<hyperv-host1-iqn>,<hyperv-host2-iqn>,<hyperv-host3-iqn>
```

2. To view the three new igroups, run the `igroup show` command.

### Map Boot LUNs to igroups

#### VMware

From the storage cluster management SSH connection, run the following commands:

```
lun map -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host1 -igroup esxi-host1 -lun-id 0
lun map -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host2 -igroup esxi-host2 -lun-id 0
lun map -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-host3 -igroup esxi-host3 -lun-id 0
lun map -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun sql-vms-datastore -igroup esxi-Hosts -lun-
id 1
lun map -vserver ESXi-Work-SVM -volume esxi_hosts_vms -lun esxi-vms-swap -igroup esxi-Hosts -lun-id 2
```

#### Hyper-V

From the storage cluster management SSH connection, run the following commands:

```
lun map -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host1 -igroup hyperv-host1 -
lun-id 0
lun map -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host2 -igroup hyperv-host2 -
lun-id 0
```

```
lun map -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun hyperv-host3 -igroup hyperv-host3 -  
lun-id 0  
lun map -vserver Hyperv-Work-SVM -volume hyperv_hosts_vms -lun sql-vms-csv -igroup hyperv-hosts -lun-  
id 1
```

# VMware vSphere 6.7U1 Setup

---

## VMware ESXi 6.7U1

This section provides detailed instructions for installing VMware ESXi 6.7U1 in an environment. After the procedures are completed, two booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

## Download Cisco Custom Image for ESXi 6.7U1

If the VMware ESXi custom image has not been downloaded, follow these steps to complete the download:

1. Click the following link: <https://my.vmware.com/web/vmware/details?downloadGroup=OEM-ESXI67U1-CISCO&productId=742>
2. You will need a user id and password on vmware.com to download this software.
3. Download the .iso file.

## Launching KVM Console Server

Cisco UCS IP KVM enables you to begin the installation of the operating system (OS) through remote media. It is necessary to log into the Cisco UCS environment to run the IP KVM. To launch the server KVM console, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. It launches the Cisco UCS Manager application. Login to the UCS manager by providing credentials.
2. Select Servers > Service Profiles > root > Sub-Organization > VM-Host-01.
3. Right-click VM-Host-01 and select KVM Console. Follow the prompts to launch the Java-based KVM console.
4. Open the KVM consoles VM-Host-02 and VM-Host-03 by following above steps from 1 to 3.

## Set Up VMware ESXi Installation

To prepare the server for the OS installation, follow these steps on each ESXi host:



**Skip this step if you are using vMedia polices. The ISO file will already be connected to KVM.**

---

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.

5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.

## Install ESXi

To install VMware ESXi to the iSCSI-bootable LUN of the hosts, follow these steps on each host:

1. Boot the server by selecting Boot Server and clicking OK. Then click OK again.
2. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, click on the Virtual Media tab and clear the  (tick) mark next to the ESXi installation media. Click Yes.



**The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.**

---

10. After the installation is complete, press Enter to reboot the server.
11. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

## Set Up Management Networking for ESXi Hosts


Adding a management network for each VMware host is necessary for managing the host.

To configure each ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as root, enter the corresponding password, and press Enter to log in.
3. Select Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.

5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.

**Figure 48 Network Adapters**



10. Press Enter.
11. Select the VLAN (Optional) option and press Enter.
12. Enter the <SQL-VM-MGMT-id> and press Enter.
13. Select IPv4 Configuration and press Enter.
14. Select the Set static IPv4 address and network configuration option by using the space bar.
15. Enter the IP address for managing the first ESXi host.
16. Enter the subnet mask for the first ESXi host.
17. Enter the default gateway for the first ESXi host.
18. Press Enter to accept the changes to the IP configuration.
19. Select the DNS Configuration option and press Enter.



---

**Because the IP address is assigned manually, the DNS information must also be entered manually.**

---

20. Enter the IP address of the primary DNS server.
21. Optional: Enter the IP address of the secondary DNS server.
22. Enter the fully qualified domain name (FQDN) for the first ESXi host.
23. Press Enter to accept the changes to the DNS configuration.
24. Press Esc to exit the Configure Management Network menu.
25. Select Test Management Network to verify that the management network is set up correctly and press Enter.
26. Press Enter to run the test, press Enter again once the test has completed, review environment if there is a failure.
27. Re-select the Configure Management Network and press Enter.
28. Select the IPv6 Configuration option and press Enter.
29. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
30. Press Esc to exit the Configure Management Network submenu.
31. Press Y to confirm the changes and reboot the ESXi host.

## Verifying UEFI secure boot of ESXi

To verify if UEFI secure boot is enabled on all the ESXi hosts, follow these steps:

1. Press ALT+F1 to open the command line console of ESXi host.
2. Log into the server with root credentials.
3. Execute `/usr/lib/vmware/secureboot/bin/secureboot.py -c`.
4. Execute `/usr/lib/vmware/secureboot/bin/secureboot.py -s`.

Figure 49 Verifying UEFI Secure Boot

```

KVM Console | Server
ESXi 6.7.0 http://www.vmware.com
Copyright (c) 2007-2018 VMware, Inc.

SQL-VM-ESX-01.aflexpod.cisco.com login: root
Password:
The time and date of this login have been sent to the system logs.

WARNING:
All commands run on the ESXi shell are logged and may be included in
support bundles. Do not provide passwords directly on the command line.
Most tools can prompt for secrets or accept them from standard input.

VMware offers supported, powerful system administration tools. Please
see www.vmware.com/go/sysadmintools for details.

The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
[root@SQL-VM-ESX-01:~] /usr/lib/vmware/secureboot/bin/secureboot.py -c
Secure boot can be enabled: All vib signatures verified. All tardisks validated. All acceptance levels validated
[root@SQL-VM-ESX-01:~] /usr/lib/vmware/secureboot/bin/secureboot.py -s
Enabled
[root@SQL-VM-ESX-01:~] _

```

## Log into VMware ESXi Hosts by Using VMware Host Client

To log into the VM-Host-01 ESXi host by using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-01 management IP address.
2. Click Open the VMware Host Client.
3. Enter root for the user name.
4. Enter the root password.
5. Click Login to connect.
6. Repeat this process to log into VM-Host-02 in a separate browser tab or window.

## Set Up Basic VMkernel Ports on Standard vSwitch

To set up the management and iSCSI VMkernel ports using standard virtual switches on the ESXi hosts, follow these steps on each host:

1. From the Host Client, select Networking.
2. In the center pane, select the Virtual switches tab.
3. Select vSwitch0.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Click Save.

7. Select Networking.
8. In the center pane, select the Virtual switches tab.
9. Select iScsiBootvSwitch.
10. Select Edit settings.
11. Change the MTU to 9000.
12. Click Save.
13. Select the VMkernel NICs tab.
14. Select vmk1 iScsiBootPG.
15. Select Edit settings.
16. Change the MTU to 9000.
17. Expand IPv4 settings and change the IP address to an address outside of the UCS SQL-FP-iSCSI-Pool-A.



**To avoid IP address conflicts if the Cisco UCS iSCSI IP Pool addresses get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.**

---

18. Click Save.
19. Select the Virtual switches tab.
20. Select the Add standard virtual switch.
21. Provide a name of iScsiBootvSwitch-B for the vSwitch Name.
22. Set the MTU to 9000.
23. Select vmnic3 from the Uplink 1 pulldown options.
24. Click Add.
25. In the center pane, select the VMkernel NICs tab.
26. Select Add VMkernel NIC
27. Specify a New port group name of iScsiBootPG-B.
28. Select iScsiBootvSwitch-B for Virtual switch.
29. Set the MTU to 9000. Do not enter a VLAN ID.
30. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask within the Configuration.



To avoid IP address conflicts, if the Cisco UCS iSCSI IP Pool addresses get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

31. Click Create.
32. On the left, select Networking, then select the Port groups tab.
33. In the center pane, right-click VM Network and select Remove.
34. Click Remove to complete removing the port group.
35. In the center pane, select Add port group.
36. Name the port group IB-MGMT Network and enter <SQL-VM-MGMT-id> in the VLAN field, and make sure Virtual Switch vSwitch0 is selected.



VMware Distributed Switch (vDS) is used for configuring remaining VMkernel ports and port groups

37. Select the VMkernel NICs tab to confirm configured virtual adapters. The adapters listed should be similar to the example shown below.

Figure 50 Basic VMkernel Configuration for Management and iSCSI Boot


Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	192.168.94.11	None
vmk1	iScsiBootPG-B	Default TCP/IP stack		192.168.12.201	None
vmk2	iScsiBootPG-B	Default TCP/IP stack		192.168.22.201	None

## Setup iSCSI Multipathing

To setup the iSCSI multipathing on the ESXi hosts, follow these steps on each ESXi host:

1. From each Host Client, select Storage.
2. In the center pane, click Adapters.
3. Select the iSCSI software adapter and click Configure iSCSI.
4. Under Dynamic targets, click Add dynamic target.
5. Enter the IP Address of iSCSI\_lif01a.
6. Repeat putting the ip address of iscsi\_lif01b, iscsi\_lif02a, iscsi\_lif02b.
7. Click Save configuration.

Figure 51 Setup ESXi iSCSI Multipath



 To get all the iscsi\_lif IP address, login to NetApp storage cluster management interface and type “network interface show” command.


 The host automatically rescans the storage adapter and the targets will be added to static targets.

## Add Required Datastores

To Add the required datastores, follow these steps on each ESXi host:

1. From the Host Client, select Storage.
2. In the center pane, select Datastores.
3. In the center pane, select New Datastore to add a new datastore.
4. In the New datastore popup, select Create new VMFS datastore and click Next.
5. Enter the name of the datastore as “SQL-VM-Datastore” and select the unclaimed device and click Next.


Figure 52 Adding Datastores



6. In the Select partitioning options page, select use full disk and VMFs 6 and click Next.
7. Click Finish to add the datastore to the ESXi host.
8. Repeat the steps (above) to add esxi-vms-swap datastore which is required to store virtual machines swap files.
9. Repeat the steps (above) on all the hosts for all the datastores (if any) that are required to be added to the ESXi hosts.

The datastore displays on the datastore list.

Figure 53 Datastores Added to ESXi Hosts




## Configure NTP on ESXi Hosts

To configure Network Time Protocol (NTP) on the ESXi hosts, follow these steps on each host:

1. From the Host Client, select Manage.
2. In the center pane, select the Time & Date tab.
3. Click Edit settings.

4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the pulldown to select Start and stop with host.
6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

**Figure 54 ESXi NTP Configuration**



7. Click Save to save the configuration changes.
8. Select Actions > NTP service > Start.
9. Verify that NTP service is now running and the clock is now set to approximately the correct time.



**The NTP server time may vary slightly from the host time.**

## VMware vCenter 6.7


The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 6.7 Server Appliance in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

### Build the VMware vCenter Server Appliance

The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, follow these steps:


1. Locate and copy the VMware-VCSA-all-6.7.0-10244745.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 6.7 vCenter Server Appliance.
2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012).
3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard displays.

Figure 55 vCenter Server Appliance 6.7 Installer




4. Click install to start the vCenter Server Appliance Deployment Wizard.
5. Click Next in the Introduction section.
6. Read and accept the license agreement and click Next.
7. In the Select deployment type section, select Embedded Platform Services Controller.
8. In the Appliance deployment target, enter the ESXi host name or IP address, User name and Password.

Figure 56 vCenter Deployment on ESXi Host



9. Click Yes to accept the certificate.
10. Enter the Appliance name and password details in the “setup appliance VM” section. Click Next.
11. In the “Select deployment size” section, select the deployment size and storage size. For example, “Tiny.”
12. Click Next.
13. In the “Select datastore” page, Select SQL-VM-Datastore and check box for Enable Think Disk Mode.
14. In the “Network Settings” section, configure the following settings:
  - a. Choose a Network: IB-MGMT Network
  - b. IP version: IPV4
  - c. IP assignment: static
  - d. System Name: <vcenter\_fqdn>
  - e. IP address: <vcenter-ip>
  - f. Subnet mask or prefix length: <vcenter-subnet-mask>
  - g. Default gateway: <vcenter-gateway>
  - h. DNS Servers: <dns-server>

Figure 57 Networking Configuration of vCenter



15. Click Next.

16. Review all values and click to finish to complete the installation.

17. The vCenter appliance installation will take few minutes to complete.

18. When installation completes, click Continue to proceed with stage 2 configuration.

19. Click Next.

20. In the Appliance Configuration, configure the following settings:

- a. Time Synchronization Mode: Synchronize time with NTP servers.
- b. NTP Servers: <ntp\_server\_ip>
- c. SSH access: Enabled.

21. Click Next.

22. Complete the SSO configuration:

- a. SSO domain name: vsphere.local
- b. SSO password: <administrator password>


- c. Confirm password: <administrator password>
- d. Site name: aflexpod
- 23. Click Next.
- 24. If needed, select join the VMware’s Customer Experience Improvement Program (CEIP).
- 25. Click Next.
- 26. Review the configuration and click Finish.
- 27. Click OK.

## Set Up VMware vCenter Server

To set up the VMware vCenter Server, follow these steps:

- 1. Using a web browser, navigate to https://<vcenter-ip>/vsphere-client.
- 2. Click Download Enhanced Authentication Plugin and install by double-clicking the downloaded file.
- 3. Login using Single Sign-On username and password created during the vCenter installation.

**Figure 58 vCenter Home**




- 4. Click Create Datacenter in the center pane.
- 5. Type AFlexPod-DC in the Datacenter name field.
- 6. Click OK.
- 7. Right-click the data center AFlexPod-DC in the list in the center pane. Click New Cluster.
- 8. Name the cluster SQL-Cluster.
- 9. Check the box to turn on the DRS. Leave the default values.
- 10. Check the box to turn on the vSphere HA. Leave the default values.

Figure 59 vCenter Cluster Creation

Name	SQL-Cluster
Location	AFlexPod-DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ——— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low ——— High
EVC	Disable
vSAN	<input type="checkbox"/> Turn ON


11. Click Ok to create the new cluster
12. On the left pane, double click the AFlexPod-DC.
13. Click Clusters.
14. Under the Clusters pane, right-click SQL-Cluster and select Settings.
15. Select Configuration > General in the list and select Edit.
16. Select Datastore specified by host and click OK.

Figure 60 Edit Cluster Settings



17. Right-click SQL-Cluster and click Add Host.
18. In the Host field, enter either the IP address or the FQDN name of one of the VMware ESXi hosts. Click Next.
19. Type root as the user name and the root password. Click Next to continue.
20. Click Yes to accept the certificate.
21. Review the host details and click Next to continue.
22. Assign a license or leave in evaluation mode and click Next to continue.
23. Click Next to continue.
24. Click Next to continue.
25. Review the configuration parameters. Click Finish to add the host.
26. Repeat the steps (above) to add all the ESXi hosts.

Figure 61 Adding ESXi Hosts to vCenter Cluster



## Configure the Default Swap File Location

To configure the default swap file location, follow these steps:

1. Expand SQL-Cluster and select ESXi host and click on configure.
2. Under virtual Machines, select Swap file location.
3. Click Edit on right pane. From the popup windows, select use a specific datastore radio button and choose esxi-vm-swap datastore.
4. Click OK

## Configure the ESXi Power Management

To configure the power management of each ESXi host, follow these steps:

1. Expand SQL-Cluster and select ESXi host and click Configure.
2. Under Hardware, select Power Management.
3. Click Edit on right pane. From the popup window, select High Performance.
4. Click OK

## Add AD User Authentication to vCenter (Optional)

To set up AD and authenticate from vCenter (If an AD Infrastructure is set up in this FlexPod environment) follow these steps:

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flexadmin (FlexPod Admin).

2. Connect to <https://<vcenter-ip>> and select Log in to vSphere Web Client.
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Navigate to Home. In the center pane, select System Configuration under Administration.
5. Select Nodes and under Nodes select the vCenter.
6. In the center pane, select the manage tab, and within the Settings select Active Directory and click Join.
7. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Click OK.
8. Right-click the vCenter and select Reboot.
9. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.
10. Log back into the vCenter Web Client.
11. In the center pane, select System Configuration under Administration.
12. Select Nodes and under Nodes select the vCenter.
13. In the center pane under the Manage tab, select Active Directory. Make sure your Active Directory Domain is listed.
14. Navigate back to the vCenter Home.
15. In the center pane under Administration, select Roles.
16. Under Single Sign-On, select Configuration.
17. Select the Identity Sources tab.
18. Click the green + sign to add an Identity Source.
19. Select the Active Directory (Integrated Windows Authentication) Identity source type.
20. Your AD domain name should be filled in. Leave Use machine account selected and click OK.
21. Your AD domain should now appear in the Identity Sources list.
22. Under Single Sign-On, select Users and Groups.
23. Select your AD domain for the Domain.
24. Make sure the FlexPod Admin user setup in step 1 appears in the list.
25. Under Administration, select Global Permissions.
26. Select the Manage tab and click the green + sign to add a User or Group.
27. In the Global Permission Root - Add Permission window, click Add.

28. In the Select Users/Groups window, select your AD Domain.

29. Under Users and Groups, select either the FlexPod Admin user or the Domain Admins group.



**The FlexPod Admin user was created in the Domain Admins group. The selection depends on whether the FlexPod Admin user will be the only user in this FlexPod or if you would like to add other users later. By selecting the Domain Admins group, any user placed in that group in the AD domain will be able to log into vCenter as an Administrator.**

---

30. Click Add. Click Check names to verify correctness of the names. Click OK to acknowledge the correctness of the names.

31. Click OK to add the selected User or Group.

32. Verify the added User or Group is listed under Users and Groups and the Administrator role is as-signed.

33. Click OK.

34. Log out and log back into the vCenter Web Client as the FlexPod Admin user. You will need to add the domain name to the user, such as flexadmin@domain.

## ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. To setup the ESXi Dump Collector, follow these steps:

1. In the vSphere web client, select Home.
2. In the center pane, click System Configuration.
3. In the left pane, select Services.
4. Under services, click VMware vSphere ESXi Dump Collector.
5. In the center pane, click the green start icon to start the service.
6. In the Actions menu, click Edit Startup Type.
7. Select Automatic.
8. Click OK.
9. Connect to each ESXi host through SSH as root.
10. Run the following commands:

```
esxcli system coredump network set -v vmk0 -j <vcenter-ip>
```

```
esxcli system coredump network set -e true
```

```
esxcli system coredump network check
```

## Cisco UCS Manager Plug-in for VMware vSphere Web Client

The Cisco UCS Manager Plug-in for VMware vSphere Web Client allows administration of Cisco UCS domains through the VMware's vCenter administrative interface. Capabilities of the plug-in include:

- View Cisco UCS physical hierarchy
- View inventory, installed firmware, faults, power and temperature statistics
- Map the ESXi host to the physical server
- Manage firmware for B and C series servers
- View VIF paths for servers
- Launch the Cisco UCS Manager GUI
- Launch the KVM consoles of UCS servers
- Switch the existing state of the locator LEDs

The installation is only valid for VMware vCenter 5.5 or higher and it requires revisions of .NET Framework 4.5 and VMware PowerCLI 5.1 or greater.

The latest UCSM plugin software can be downloadable from:

<https://software.cisco.com/download/home/286282669/type>

For the user guides and installation guides, go to: <https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html>

## FlexPod VMware vSphere Distributed Switch (vDS)

This section provides the procedures for installing the VMware vDS on the FlexPod ESXi Management Hosts.

In the Cisco UCS setup section of this document two sets of vNICs (Host-A and B, and iSCSI-A and B) were setup. The vmnic ports associated with the Host-A and B vNICs will be migrated to VMware vDS in this procedure. The critical infrastructure VLAN interfaces and vMotion interfaces will be configured on the vDS.

An IB-Mgmt VLAN and a SQL-VM-Traffic VLAN port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN cloud, to the Cisco UCS Host-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC and peer-link interfaces on the switches.


### Configure the VMware vDS in vCenter

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere Web Client, select Networking under the Home tab.
2. Right-click the AFlexPod-DC datacenter and select Distributed Switch > New Distributed Switch.
3. Give the Distributed Switch a descriptive name and click Next.
4. Make sure Distributed switch: 6.6.0 is selected and click Next.


5. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control Enabled. Otherwise, Disable Network I/O Control. Enter SQL-VM-Traffic for the Port group name. Click Next.
6. Review the information and click Finish to complete creating the vDS.

**Figure 62 Creating Distributed Switch**




7. Expand the AFlexPod-DC datacenter and the newly created vDS. Select the newly created vDS.
8. Select the vDS. Click the Edit distributed switch settings icon.
9. Select the vDS. Click the Edit distributed switch settings icon.
10. On the left in the Edit Settings window, select Advanced.
11. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

Figure 63 Set MTU on vDS to 9000




12. Select the SQL-VM-Traffic port group. In the center pane, select the Edit distributed port group settings icon. The Edit button can be used to change the number of ports in the port group to a number larger than the default of 8. All the other properties of the port group can also be changed under Edit.
13. For Advanced, Security, Traffic Shaping tabs leave values at defaults.
14. On the VLAN tab, Choose VLAN type as VLAN and set the VLAN ID as <904> or <IB-Mgmt-VLAN>
15. On the Teaming and failover tab, move both uplink 1 and 2 under active uplinks.

Figure 64 Uplink Ports for SQL Virtual Machine Traffic



16. For Monitoring, traffic filtering and marking, miscellaneous tabs leave values at defaults.
17. Click Save.
18. Multiple Distributed Port groups will be created for service different network traffics. First, create port group for VMkernel adapter.
19. Right-click the vDS, select Distributed Port Group > New Distributed Port Group.
20. Enter the name as IB-MGMT and click Next.
21. Change the VLAN Type as VLAN and enter the VLAN ID for IB-MGMT traffic.
22. Check the Customize default policies configuration check box and click Next.
23. Leave the defaults values for Security and Traffic shaping tabs and click Next.
24. On the Teaming and failover tab, move uplink 1 under Active Uplink and uplink 2 under Standby uplinks as shown below.


**Figure 65 Uplink Configuration for IB-MGMT Traffic**



25. For the Monitoring, Miscellaneous and Edit additional settings tabs, leave the settings defaults and click Next.
26. Review the changes once and then click Finish to create the port group for IB-MGMT traffic.
27. Create the second port group for vMotion VMkernel for vMotion traffic. Right-click the vDS and select New Distributed Port Group.
28. Enter the name as SQL-vMotion and click Next.
29. Change the VLAN Type as VLAN and enter the VLAN ID for vMotion traffic.


30. Check the Customize default policies configuration check box and click Next.
31. Leave the defaults values for Security and Traffic shaping tabs and click Next.
32. On the Teaming and failover tab, move uplink 2 under Active Uplink and uplink 1 under Standby uplinks as shown below.

**Figure 66 Uplink Configuration for vMotion Port Group**




33. For the Monitoring, Miscellaneous and Edit additional settings tabs, leave the settings defaults and click Next.
34. Review the changes once and then click Finish to create the port group for SQL-vMotion traffic.
35. Create Port Groups for SQL guest to use for accessing NetApp storage volumes using In-Guest iSCSI software initiator. Right-click the vDS, select Distributed Port Group > New Distributed Port Group.
36. Enter the name as SQL-Guest-iSCSI-A and click Next.
37. Change the VLAN Type as VLAN and enter 3012 as VLAN ID for storage traffic passing through Fabric A.
38. Check the Customize default policies configuration check box and click Next.
39. Leave the defaults values for Security and Traffic shaping tabs and click Next.
40. On the Teaming and failover tab, move uplink 1 under Active Uplink and uplink 2 under Standby uplinks as shown below.

Figure 67 Uplinks Configuration SQL-Guest-iSCSI-A Port Group



41. For the Monitoring, Miscellaneous and Edit additional settings tabs, leave the settings defaults and click Next.
42. Review the changes once and then click Finish to create the port group for SQL-Guest-iSCSI-A traffic.
43. Create another port group for In-Guest storage access using Fabric B path. Right-click the vDS, select Distributed Port Group > New Distributed Port Group.
44. Enter the name as SQL-Guest-iSCSI-B and click Next.
45. Change the VLAN Type as VLAN and enter 3022 as VLAN ID for storage traffic passing through Fabric A.
46. Check the Customize default policies configuration check box and click Next.
47. Leave the defaults values for Security and Traffic shaping tabs and click Next.
48. On the Teaming and failover tab, move uplink 2 under Active Uplink and uplink 1 under Standby uplinks as shown below.

Figure 68 Uplinks Configuration SQL-Guest-iSCSI-B Port Group



49. For the Monitoring, Miscellaneous and Edit additional settings tabs, leave the settings defaults and click Next.


50. Review the changes once and then click Finish to create the port group for SQL-Guest-iSCSI-B traffic.

51. To add the ESXi nodes to the vDS, right-click vDS and select Add and Manage Hosts.

52. Make sure Add hosts is selected and click Next.


53. Click the green + sign to add hosts. Select the ESXi hosts and click OK. Select Configure identical network settings on multiple hosts check box and Click Next.

Figure 69 Adding ESXi Hosts to vDS Migration



- 54. Since all hosts have the same configuration, you can select any host as template host. Click Next.
- 55. Leave Manage physical adapters and Manage VMkernel adapters selected. Select Migrate virtual machine networking and click Next.
- 56. Select vmnic0 on the template server and click on Assign uplink. Select Uplink 1 and click OK.
- 57. Select vmnic1 on the template server and click on Assign uplink. Select Uplink 2 and click OK.
- 58. Click Apply to all to assign vmnic0 to uplink 1 and vmnic1 to uplink 2 on all the hosts as shown below.

**Figure 70 Managing Physical Adapters in vDS**




- 59. Select vmk0 on the template server and click on Assign port group. Select the IB-MGMT destination port group and click OK.
- 60. Click Apply to all to apply the vmk0 port to IB-MGMT on all the hosts.

 A warning message about the loss of networking connectivity during the migration activity appears. It can be ignored at this moment since no workload virtual machines are deployed yet.


- 61. On the popup windows, provide the management IP addresses of the other two hosts as shown below.

Figure 71 vDS VMkernel Migration



62. Click OK to migrate vmk0 to IB-MGMT on all the hosts as shown below.

Figure 72 Migrating VMkernel Network Adapters



63. Click Next.




78. Select VMkernel Network Adapter and click Next.

79. On the select target device window, click Select an existing network and browse to SQL-vMotion port group. Click Next.

80. On the Port properties window, select vMotion for TCP/IP Stack and click Next.

**Figure 74 Adding VMkernel Adapter Using vDS**



81. On IPV4 settings window, select use static IPV4 settings radio button and provide IPV4 address and subnet mask. Click Next.

82. Review the details and click Finish to create VMkernel port for vMotion.


83. On the middle pane, select the VMkernel adapters and select vmk3 adapter and click on Edit settings.

84. Under NIC settings, set MTU as 9000 and click OK.

85. Repeat steps from 76 to 84 to add VMkernel Adapter on other two hosts.

86. On each host, the vDS configuration should look similar as shown below.

Figure 75 Final vDS Configuration




## Creating Virtual Machine for SQL Server

To create a virtual machine for SQL Server database application deployment, follow these steps:

1. Log into vCenter web client using vCenter IP address.
2. Expand AFlexPod-DC. Right click on SQL-Cluster cluster and select New Virtual machine.
3. On the New Virtual Machine window, Select Create a new virtual machine and click Next.
4. Type name as SQLVM01 and select AFlexPod-DC as location for the virtual machine and click Next.
5. For Compute resource, expand the SQL-Cluster and select the any ESXi host. Click Next.
6. Select SQL-VM-Datastore for storage. Click Next.
7. Choose ESXi6.7 and later for virtual machine compatibility and click Next.
8. Choose Linux for OS Gest Family and Red Hat Enterprise Linux 7 (64-bit) as Guest OS Version. Click Next.
9. Under Virtual Hardware tab, set the hardware resource appropriately as follows:
  - a. Set CPU to 4.
  - b. Set Memory to 16GB (16,384MB) and select Reserve all guest memory option.
  - c. Set New Hard disk size to 40GB and choose SQL-VM-Datastore for Location.
  - d. Set Disk Provisioning Thick provision lazy zeroed.
  - e. Set New Network to "SQL-VM-Traffic(SQL-vDS)" port group. Make sure Connect At Power On is selected and Adapter Type is set to VMNEXT3.


- f. Click New device drop-down list and select Network. Add two additional network adaptors for In-Guest iSCSI storage configuration.
- g. On the first new network adaptor, select “SQL-Guest-iSCSI-A(SQL-vDS)” and for the second new network adapter select “SQL-Guest-iSCSI-B(SQL-vDS)”.
- h. Browse through the SQL-VM-Datastore and select RHEL7.4 iso file for New CD/DVD Drive. Select Connect At Power On option.

Figure 76 Creating SQL Virtual Machine on VMware Cluster



- 10. Select VM Options, expand Boot Options and set Firmware to EFI and select the Secure Boot check box. (optional).

Figure 77 Virtual Machine UEFI Secure Boot



11. Click Next and review the details.
12. Click Finish to create the virtual Machine.
13. When the virtual machine is created, right-click then new virtual machine and select Power On.
14. To Open the console of the virtual machine, right-click the virtual machine and select Open Console.
15. Repeat the steps to create minimum of three virtual machines running RHEL7.4 OS.

# Microsoft Windows Server 2016 Hyper-V Deployment

---

## Install Microsoft Windows Server 2016

This section provides the instructions to install Microsoft Windows Server 2016 on Cisco UCS B200 M5 servers.

Several methods exist for installing Microsoft Windows Server 2016. The provided steps focus on how to use the built-in keyboard, video, mouse (KVM) console and leverage the virtual media features of Cisco UCS Manager. UCSM maps the remote installation media to individual servers and connects to their boot LUNs.

The Cisco UCS IP KVM enables you to begin the installation of the operating system (OS) through remote media. It is necessary to log into the UCS environment to run the IP KVM.

To log into the Cisco UCS environment, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log into Cisco UCS Manager, click Login.
6. From the main menu, click Servers.
7. Select Servers > Service Profiles > root > Sub-Organization > VM-Host-01
8. Right-click VM-Host-01 and select KVM Console.
9. Follow the prompts to launch the Java-based KVM console.
10. Select Servers > Service Profiles > Sub-Organization > VM-Host-01.
11. Right-click VM-Host-01. and select KVM Console.
12. Follow the prompts to launch the Java-based KVM console.
13. From the virtual KVM Console, select the Virtual Media tab.
14. Select Add Image.
15. Browse to the Windows Server 2016 installation ISO image file and click Open.
16. Map the image that you just added by selecting Mapped.
17. To boot the server, select the KVM tab.
18. Select Power On Server in the KVM interface Summary tab, and then click OK.

## Install Windows Server 2016

To install Windows Server 2016 on each host, follow these steps:

1. On boot, the machine detects the presence of the Windows installation media.
2. After the installer has finished loading, Enter the relevant region information and click Next.
3. Click Install now.
4. Enter the Product Key and click Next.
5. Select Windows Server 2016 Datacenter (Desktop Experience) and click Next.




**Optionally, you may remove the GUI after the Hyper-V cluster is operational.**

---


6. After reviewing the EULA, accept the license terms and click Next.
7. Select Custom: Install Windows only (advanced).
8. Select Custom (advanced) installation.
9. In the Virtual Media Session manager uncheck the Mapped checkbox for the Windows ISO and select yes to confirm.
10. Click Add Image.
11. Browse to the Cisco fNIC driver ISO, click Open.
12. Check the Mapped checkbox next to the Cisco eNIC Driver ISO. Download the latest driver iso image from the cisco.com site.

**Figure 78** Installing eNIC Drivers for Windows



13. Back in the KVM Console, click Load Driver and then, click OK.
14. The Cisco VIC FCoE Storport Miniport driver should auto detected, click Next.

Figure 79 eNIC Driver for Windows



15. You will see a LUN listed in the drive selection screen.

---

 Only single LUN instance should be displayed. If multiple instances of the same LUN are seen, it indicates that there are multiple paths to the boot LUN which will cause installation issues since there are no multipath drivers available/enabled at this stage.

---




---

 The message "Windows Can't be installed on this drive" appears because the Windows installation ISO image is not mapped at this time.

---

Figure 80 Disk During Windows Installation



16. Select the LUN and click Next to continue with the install.


17. At the end of the installation, enter an administrator password on the settings page and click Finish.

### Install Chipset Driver Software

This section provides detailed information on installing the Intel chipset driver software. To install the Chipset drivers on each Hyper-V host, follow these steps:

1. In the Virtual Media Session manager, uncheck the Mapped checkbox for the Windows ISO.
2. Click Add Image.
3. Browse to the Cisco UCS driver ISO, click Open.
4. Check the Mapped checkbox for the Cisco UCS driver ISO.
5. Browse to the CD ROM > Chipset > Intel > <Server Model> W2K16 > x64.

**Figure 81 Installation Chipset Drivers for Windows**



6. Double-click Setup Chipset to install the chipset driver and reboot the system.

## Install Windows Roles and Features

This section explains how to add/install the MPIO, Failover Clustering and Hyper-V roles and feature using PowerShell. To install Windows Roles and Features all the Hyper-V nodes, follow these steps:

1. Open PowerShell window (with Run As Administrator) and run the following command to install Windows server roles and features.

```
Install-WindowsFeature -Name 'Multipath-I/O','Failover-Clustering','Hyper-V' -IncludeManagementTools
```

**Figure 82 Installing Windows Roles and Features**

```
PS C:\Windows\system32> Get-WindowsFeature -Name 'Multipath-I/O','Failover-Clustering','Hyper-V'
```

Display Name	Name	Install State
[ ] Hyper-V	Hyper-V	Available
[ ] Failover Clustering	Failover-Clustering	Available
[ ] Multipath I/O	Multipath-I/O	Available

```
PS C:\Windows\system32> Install-WindowsFeature -Name 'Multipath-I/O','Failover-Clustering','Hyper-V' -IncludeManagementTools
```

2. Restart the computer and verify the roles and features are installed as shown below.

Figure 83 Verifying Windows Roles and Features

```
PS C:\Windows\system32> Get-WindowsFeature -Name 'Multipath-IO','Failover-Clustering','Hyper-V'
```

Display Name	Name	Install State
[X] Hyper-V	Hyper-V	Installed
[X] Failover Clustering	Failover-Clustering	Installed
[X] Multipath I/O	Multipath-IO	Installed

## Configure Networking

To configure networking, follow these steps:

1. Gather the network adapter details and create the SET virtual switch using the following PowerShell cmdlets:

```
Get-NetAdapter
```

```
New-VMSwitch -Name SETswitch -NetAdapterName "00-Host-A","01-Host-B" -
EnableEmbeddedTeaming $true
```

Figure 84 Creating SET Switch

```
PS C:\Users\Administrator> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
03-iSCSI-B 2	Cisco VIC Ethernet Interface #2	3	Up	00-25-B5-B9-08-0D	40 Gbps
01-Host-B	Cisco VIC Ethernet Interface #4	4	Up	00-25-B5-B9-08-0C	40 Gbps
02-iSCSI-A	Cisco VIC Ethernet Interface	10	Up	00-25-B5-B9-0A-0D	40 Gbps
00-Host-A	Cisco VIC Ethernet Interface #3	5	Up	00-25-B5-B9-0A-0C	40 Gbps

```
PS C:\Users\Administrator> New-VMSwitch -Name SETswitch -NetAdapterName "00-Host-A","01-Host-B" -EnableEmbeddedTeaming $true
```

Name	SwitchType	NetAdapter	InterfaceDescription
SETswitch	External	Teamed-Interface	

2. Create the virtual network adapters on the management OS and connect them to the virtual switch using the following PowerShell cmdlets:

```
Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "MGMT" -ManagementOS
```

```
Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "LM" -ManagementOS
```

```
Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "Cluster" -ManagementOS
```

Figure 85 Creating Parent Node Network Interfaces

```
PS C:\Users\Administrator> Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "MGMT" -ManagementOS
PS C:\Users\Administrator> Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "LM" -ManagementOS
PS C:\Users\Administrator> Add-VMNetworkAdapter -SwitchName "SETswitch" -Name "Cluster" -ManagementOS
```

3. Assign VLAN ID to the virtual network adapters created in the previous step:

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName MGMT -ManagementOS -Access -VlanId 906
```

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName LM -ManagementOS -Access -VlanId 907
```

```
Set-VMNetworkAdapterVlan -VMNetworkAdapterName Cluster -ManagementOS -Access -VlanId 908
```

Figure 86 Setting VLANs On Network Interfaces

```
PS C:\Users\Administrator> Set-VMNetworkAdapterVlan -VMNetworkAdapterName MGMT -ManagementOS -Access -VlanId 906
PS C:\Users\Administrator> Set-VMNetworkAdapterVlan -VMNetworkAdapterName LM -ManagementOS -Access -VlanId 907
PS C:\Users\Administrator> Set-VMNetworkAdapterVlan -VMNetworkAdapterName Cluster -ManagementOS -Access -VlanId 908
```

- Verify the virtual network adapters.

Figure 87 Listing Network Interfaces

```
PS C:\Users\Administrator> Get-NetAdapter | ft -AutoSize
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
vEthernet (Cluster)	Hyper-V Virtual Ethernet Adapter #4	18	Up	00-15-5D-0D-6D-02	80 Gbps
vEthernet (LM)	Hyper-V Virtual Ethernet Adapter #3	5	Up	00-15-5D-0D-6D-01	80 Gbps
vEthernet (MGMT)	Hyper-V Virtual Ethernet Adapter #2	4	Up	00-15-5D-0D-6D-00	80 Gbps
vEthernet (SETswitch)	Hyper-V Virtual Ethernet Adapter	17	Up	00-25-B5-B9-0A-0A	80 Gbps
03-iSCSI-B 2	Cisco VIC Ethernet Interface #2	15	Up	00-25-B5-B9-0B-0B	40 Gbps
01-Host-B	Cisco VIC Ethernet Interface #4	2	Up	00-25-B5-B9-0B-0A	40 Gbps
02-iSCSI-A	Cisco VIC Ethernet Interface	14	Up	00-25-B5-B9-0A-0B	40 Gbps
00-Host-A	Cisco VIC Ethernet Interface #3	3	Up	00-25-B5-B9-0A-0A	40 Gbps

- Assign static IP address to the "MGMT" virtual network adapter and configure the DNS server address.

```
New-NetIPAddress -ifIndex 4 -IPAddress 192.168.96.109 -PrefixLength 24 -
DefaultGateway 192.168.96.254
```

Figure 88 Assigning IP Addresses

```
PS C:\Users\Administrator> New-NetIPAddress -ifIndex 19 -IPAddress 192.168.96.11 -PrefixLength 24 -DefaultGateway 192.168.96.254
```

```
IPAddress      : 192.168.96.11
InterfaceIndex : 19
InterfaceAlias : vEthernet (MGMT)
AddressFamily  : IPv4
Type           : Unicast
PrefixLength   : 24
PrefixOrigin   : Manual
SuffixOrigin   : Manual
AddressState   : Tentative
ValidLifetime  : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)
SkipAsSource   : False
PolicyStore    : ActiveStore
```

```
IPAddress      : 192.168.96.11
InterfaceIndex : 19
InterfaceAlias : vEthernet (MGMT)
AddressFamily  : IPv4
Type           : Unicast
PrefixLength   : 24
PrefixOrigin   : Manual
SuffixOrigin   : Manual
AddressState   : Invalid
ValidLifetime  : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)
SkipAsSource   : False
PolicyStore    : PersistentStore
```

```
Set-DnsClientServerAddress -InterfaceAlias "vEthernet (MGMT)" -ServerAddresses
10.1.166.9
```

Figure 89 Setting DNS IP Address

```
PS C:\Users\Administrator> Set-DnsClientServerAddress -InterfaceAlias "vEthernet (MGMT)" -ServerAddresses 10.1.166.9
```

- Assign static IP addresses to the "LM" virtual network adapter.

```
New-NetIPAddress -ifAlias "vEthernet (LM)" -IPAddress 192.168.97.109 -PrefixLength
24
```

- Assign static IP addresses to the "Cluster" virtual network adapter.

```
New-NetIPAddress -ifAlias "vEthernet (Cluster)" -IPAddress 192.168.98.109 -
PrefixLength 24
```

8. Disable IPV6 on all Network adapters (optional).

```
Get-NetAdapterBinding -infa 2
```

Figure 90 Configuring Network Binding

```
PS C:\Windows\system32> Get-NetAdapterBinding -infa 2
```

Name	DisplayName	ComponentID	Enabled
vEthernet (Cluster)	Microsoft LLDP Protocol Driver	ms_lldp	True
vEthernet (Cluster)	Link-Layer Topology Discovery Mapper I/O Driver	ms_lltdio	True
vEthernet (Cluster)	Microsoft Network Adapter Multiplexor Protocol	ms_implat	False
vEthernet (Cluster)	Internet Protocol Version 4 (TCP/IPv4)	ms_tcpip	True
vEthernet (Cluster)	File and Printer Sharing for Microsoft Networks	ms_server	True
vEthernet (Cluster)	Client for Microsoft Networks	ms_msclient	True
vEthernet (Cluster)	Link-Layer Topology Discovery Responder	ms_rspndr	True
vEthernet (Cluster)	Internet Protocol Version 6 (TCP/IPv6)	ms_tcpip6	True
vEthernet (Cluster)	Hyper-V Extensible Virtual Switch	vms_pp	False
vEthernet (Cluster)	QoS Packet Scheduler	ms_pacer	True
vEthernet (LM)	Client for Microsoft Networks	ms_msclient	True
vEthernet (LM)	Link-Layer Topology Discovery Responder	ms_rspndr	True
vEthernet (LM)	QoS Packet Scheduler	ms_pacer	True
vEthernet (LM)	Microsoft LLDP Protocol Driver	ms_lldp	True
vEthernet (LM)	Link-Layer Topology Discovery Mapper I/O Driver	ms_lltdio	True
vEthernet (LM)	Microsoft Network Adapter Multiplexor Protocol	ms_implat	False
vEthernet (LM)	Internet Protocol Version 4 (TCP/IPv4)	ms_tcpip	True
vEthernet (LM)	Hyper-V Extensible Virtual Switch	vms_pp	False
vEthernet (LM)	Internet Protocol Version 6 (TCP/IPv6)	ms_tcpip6	True

```
Disable-NetAdapterBinding -infa 2 -ComponentID ms_tcpip6 -Name vE*
```

```
Disable-NetAdapterBinding -infa 2 -ComponentID ms_tcpip6 -Name 0*
```

Figure 91 Disabling Network Binding

```
PS C:\Windows\system32> Disable-NetAdapterBinding -infa 2 -ComponentID ms_tcpip6 -Name vE*
PS C:\Windows\system32> Disable-NetAdapterBinding -infa 2 -ComponentID ms_tcpip6 -Name 0*
```

9. Enable jumbo/mtu on interfaces used for Live Migration traffic.

```
Get-NetAdapterAdvancedProperty -Name "vEthernet (LM)"
```

```
Set-NetAdapterAdvancedProperty -Name "vEthernet (LM)" -DisplayName "Jumbo Packet" -
DisplayValue "9014 Bytes"
```

Figure 92 Configuring MTU

```
PS C:\Windows\system32> Set-NetAdapterAdvancedProperty -Name "vE* (LM)" -DisplayName "Jumbo Packet" -DisplayValue "9014 Bytes"
PS C:\Windows\system32> Get-NetAdapterAdvancedProperty -Name "vE* (LM)"
```

Name	DisplayName	DisplayValue	RegistryKeyword	RegistryValue
vEthernet (LM)	IPv4 Checksum Offload	Rx & Tx Enabled	*IPChecksum0...	{3}
vEthernet (LM)	IPSec Offload	Auth Header and ESP Enabled	*IPsecOffloadV2	{3}
vEthernet (LM)	Jumbo Packet	9014 Bytes	*JumboPacket	{9014}
vEthernet (LM)	Large Send Offload Version ...	Enabled	*LsoV2IPv4	{1}
vEthernet (LM)	Large Send Offload Version ...	Enabled	*LsoV2IPv6	{1}
vEthernet (LM)	Network Direct (RDMA)	Disabled	*NetworkDirect	{0}
vEthernet (LM)	Receive Side Scaling	Enabled	*RSS	{1}
vEthernet (LM)	TCP Checksum Offload (IPv4)	Rx & Tx Enabled	*TCPChecksum...	{3}
vEthernet (LM)	TCP Checksum Offload (IPv6)	Rx & Tx Enabled	*TCPChecksum...	{3}
vEthernet (LM)	UDP Checksum Offload (IPv4)	Rx & Tx Enabled	*UDPChecksum...	{3}
vEthernet (LM)	UDP Checksum Offload (IPv6)	Rx & Tx Enabled	*UDPChecksum...	{3}
vEthernet (LM)	Network Address	--	NetworkAddress	{--}

### Install NetApp Host Utilities


After enabling the MPIO feature in Windows, download and install NetApp Windows Unified Host Utilities on all the Hyper-V nodes. To download and install the host utilities, follow these steps:

1. Download the NetApp host utilities v7.1 for Windows from the link below:

<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=61343>

2. Unzip the file and run the executable file. The NetApp Windows Unified Host Utilities setup wizard is launched and click Next.
3. Select “Yes, install support for Multipath IO” and click Next.

Figure 93 NetApp Host Utilities for Windows



4. Accept the default destination folder and click Next.
5. Click Next and Install to start the installation of host utilities.
6. After the installation is complete, click Finish and restart the computer.

## Configure Multipath-IO

To configure Multipath-IO, follow these steps:

1. In the previous section, the MPIO feature is enabled. If it is not enabled, then run the following PowerShell command to enable it:

```
Install-WindowsFeature -Name Multipath-IO
```

Figure 94 Installing Multipath Drivers

```
PS C:\Users\Administrator> Install-WindowsFeature -Name Multipath-IO
```

2. Enable Microsoft Device Specific Module (MSDSM) to automatically claim SAN disks for Microsoft Multipath I/O (MPIO) for iSCSI bus type using the following PowerShell cmdlet:


```
Enable-MSDSMAutomaticClaim -BusType iSCSI
```

Figure 95 Enabling Device Modules

```
PS C:\Users\Administrator> Enable-MSDSMAutomaticClaim -BusType iSCSI
```

3. Restart the computer.
4. Verify the MPIO by navigating to Server Manager > Tools > MPIO.

Figure 96 Adding NetApp Device IDs



5. Go back to section [Configure Server Boot Order](#) and follow steps 12 through 30 to add the remaining paths/targets.


## Host Renaming and Join to Domain

To rename the host, follow these steps:

1. Log into the host and open PowerShell.
2. Rename the host:

```
Rename-Computer -NewName <hostname> -restart
```

Figure 97 Renaming Host



3. Add the host to Active Directory domain as shown below.

Figure 98 Joining the Host to Domain

```
PS C:\Users\Administrator> Add-Computer -DomainName aflexpod.cisco.com

cmdlet Add-Computer at command pipeline position 1
Supply values for the following parameters:
Credential
WARNING: The changes will take effect after you restart the computer FP-HV-Host1.
PS C:\Users\Administrator>
```

### Configure Microsoft iSCSI Initiator

To configure the Microsoft iSCSI initiator, follow these steps:

1. Enable the iSCSI initiator by using Windows PowerShell to start the Microsoft iSCSI Service and change the service startup type to Automatic by running these two commands:

```
Start-Service msiscsi
```

Figure 99 Starting iSCSI initiator

```
PS C:\Windows\system32> Start-Service -Name MSiSCSI
```


```
Set-Service msiscsi -startuptype "automatic"
```

Figure 100 Enabling iSCSI Initiator

```
PS C:\Windows\system32> Set-Service -Name MSiSCSI -StartupType Automatic
```

2. Click the Targets tab and you will see a discovered target with status as connected.

Figure 101 iSCSI Initiator Configuration



- Establish a connection to a new iSCSI Target Portal using the following PowerShell cmdlet.

```
New-IscsiTargetPortal -TargetPortalAddress 192.168.13.18 -InitiatorPortalAddress 192.168.13.108
```

```
New-IscsiTargetPortal -TargetPortalAddress 192.168.13.19 -InitiatorPortalAddress 192.168.13.108
```

```
New-IscsiTargetPortal -TargetPortalAddress 192.168.23.18 -InitiatorPortalAddress 192.168.23.108
```

```
New-IscsiTargetPortal -TargetPortalAddress 192.168.23.19 -InitiatorPortalAddress 192.168.23.108
```

Figure 102 Connecting to NetApp Storage Using iSCSI Initiator

```
PS C:\Windows\system32> New-IscsiTargetPortal -TargetPortalAddress 192.168.13.18 -InitiatorPortalAddress 192.168.13.109

InitiatorInstanceName : ROOT\ISCSIPRT\0000_0
InitiatorPortalAddress : 192.168.13.109
IsDataDigest          : False
IsHeaderDigest        : False
TargetPortalAddress   : 192.168.13.18
TargetPortalPortNumber : 3260
PSComputerName        :

PS C:\Windows\system32> New-IscsiTargetPortal -TargetPortalAddress 192.168.13.19 -InitiatorPortalAddress 192.168.13.109

InitiatorInstanceName : ROOT\ISCSIPRT\0000_0
InitiatorPortalAddress : 192.168.13.109
IsDataDigest          : False
IsHeaderDigest        : False
TargetPortalAddress   : 192.168.13.19
TargetPortalPortNumber : 3260
PSComputerName        :

PS C:\Windows\system32> New-IscsiTargetPortal -TargetPortalAddress 192.168.23.18 -InitiatorPortalAddress 192.168.23.109

InitiatorInstanceName : ROOT\ISCSIPRT\0000_0
InitiatorPortalAddress : 192.168.23.109
IsDataDigest          : False
IsHeaderDigest        : False
TargetPortalAddress   : 192.168.23.18
TargetPortalPortNumber : 3260
PSComputerName        :

PS C:\Windows\system32> New-IscsiTargetPortal -TargetPortalAddress 192.168.23.19 -InitiatorPortalAddress 192.168.23.109

InitiatorInstanceName : ROOT\ISCSIPRT\0000_0
InitiatorPortalAddress : 192.168.23.109
IsDataDigest          : False
```

- After a new target portal is created, the iSCSI target, and its connection status are displayed through the Get-IscsiTarget cmdlet as shown below:

Figure 103 Listing iSCSI Targets

```
PS C:\Windows\system32> Get-IscsiTarget

IsConnected NodeAddress                                     PSComputerName
-----
True iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7
```

5. Connect to all available iSCSI Targets, use the following command:

```
$target = Get-IscsiTarget
```

```
Connect-IscsiTarget -TargetPortalAddress 192.168.13.18 -InitiatorPortalAddress
192.168.13.108 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -
IsPersistent $true
```

Figure 104 Adding First iSCSI Target

```
PS C:\Windows\system32> $target = Get-IscsiTarget

PS C:\Windows\system32> $target.NodeAddress
iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7

PS C:\Windows\system32> Connect-IscsiTarget -TargetPortalAddress 192.168.13.18 -InitiatorPortalAddress 192.168.13.109 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -IsPersistent $true

AuthenticationType : NONE
InitiatorInstanceName : ROOT\iScsiPrt\0000_0
InitiatorNodeAddress : iqn.2010-11.com.flexpod:bb09-6332-host:6
InitiatorPortalAddress : 192.168.13.109
InitiatorSideIdentifier : 400001370002
IsConnected : True
IsDataDigest : False
IsDiscovered : True
IsHeaderDigest : False
IsPersistent : True
NumberOfConnections : 1
SessionIdentifier : fffffa4813152f010-4000013700000007
TargetNodeAddress : iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7
TargetSideIdentifier : 1400
PSComputerName :
```

```
Connect-IscsiTarget -TargetPortalAddress 192.168.13.19 -InitiatorPortalAddress
192.168.13.108 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -
IsPersistent $true
```

```
Connect-IscsiTarget -TargetPortalAddress 192.168.23.18 -InitiatorPortalAddress
192.168.23.108 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -
IsPersistent $true
```

Figure 105 Adding Second and Third iSCSI Target

```
PS C:\Windows\system32> Connect-IscsiTarget -TargetPortalAddress 192.168.13.19 -InitiatorPortalAddress 192.168.13.109 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -IsPersistent $true

AuthenticationType : NONE
InitiatorInstanceName : ROOT\iScsiPrt\0000_0
InitiatorNodeAddress : iqn.2010-11.com.flexpod:bb09-6332-host:6
InitiatorPortalAddress : 192.168.13.109
InitiatorSideIdentifier : 400001370004
IsConnected : True
IsDataDigest : False
IsDiscovered : True
IsHeaderDigest : False
IsPersistent : True
NumberOfConnections : 1
SessionIdentifier : fffffa4813152f010-4000013700000008
TargetNodeAddress : iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7
TargetSideIdentifier : 1400
PSComputerName :
```

```
PS C:\Windows\system32> Connect-IscsiTarget -TargetPortalAddress 192.168.23.18 -InitiatorPortalAddress 192.168.23.109 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -IsPersistent $true

AuthenticationType : NONE
InitiatorInstanceName : ROOT\iScsiPrt\0000_0
InitiatorNodeAddress : iqn.2010-11.com.flexpod:bb09-6332-host:6
InitiatorPortalAddress : 192.168.23.109
InitiatorSideIdentifier : 400001370006
IsConnected : True
IsDataDigest : False
IsDiscovered : True
IsHeaderDigest : False
IsPersistent : True
NumberOfConnections : 1
SessionIdentifier : fffffa4813152f010-4000013700000009
TargetNodeAddress : iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7
TargetSideIdentifier : 1900
PSComputerName :
```

```
Connect-IscsiTarget -TargetPortalAddress 192.168.23.19 -InitiatorPortalAddress
192.168.23.108 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -
IsPersistent $true
```

Figure 106 Adding Fourth iSCSI Target

```
PS C:\Windows\system32> Connect-IscsiTarget -TargetPortalAddress 192.168.23.19 -InitiatorPortalAddress 192.168.23.109 -NodeAddress $target.NodeAddress -IsMultipathEnabled $true -IsPersistent $true

AuthenticationType      : NONE
InitiatorInstanceName   : ROOT\iScsiPrt\0000_0
InitiatorNodeAddress    : iqn.2010-11.com.flexpod:bb09-6332-host:6
InitiatorPortalAddress  : 192.168.23.109
InitiatorSideIdentifier : 400001370008
IsConnected             : True
IsDataDigest           : False
IsDiscovered            : True
IsHeaderDigest         : False
IsPersistent            : True
NumberOfConnections     : 1
SessionIdentifier       : fffffa4813152f010-4000013700000000a
TargetNodeAddress       : iqn.1992-08.com.netapp:sn.e602f9c3f8b211e892c200a098a9fed2:vs.7
TargetSideIdentifier    : 1500
PSComputerName         :
```

6. Verify the Discovery and Favorite Targets in the iSCSI initiator properties page as shown below .

Figure 107 Adding all the iSCSI Targets






Figure 108 Verifying iSCSI Target Paths



7. After the above steps and configuration on the iSCSI storage array, the iSCSI disks should be visible in the Computer Management > Disk Management as shown below.

Figure 109 Listing Disk using Disk Management



## Create Windows Failover Cluster

This section explains how to create a failover cluster by using Windows PowerShell. Make sure the Failover Cluster feature is installed on all the Hyper-V hosts.


### Prepare the Shared Disks for Cluster Use

To prepare the shared disks for cluster use, follow these steps:

1. Log into one of the Hyper-V hosts and run the following PowerShell command to gather information about the disks for cluster use. Write down the Number and Total Size of the shared disks from the output.

```
Get-Disk | where PartitionStyle -EQ 'RAW'
```

Figure 110 Listing Disks



- On the same Hyper-V hosts, run the following PowerShell command to prepare the disk for cluster quorum by initializing, creating a new partition, and formatting it.

```
get-disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 1 | `
Initialize-Disk -PartitionStyle MBR -PassThru | `
New-Partition -AssignDriveLetter -UseMaximumSize | `
Format-Volume -FileSystem NTFS -NewFileSystemLabel "Quorum" -Confirm:$false
```

Figure 111 Listing Quorum Disk

```
PS C:\Windows\system32> get-disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 1 | `
Initialize-Disk -PartitionStyle MBR -PassThru | `
New-Partition -AssignDriveLetter -UseMaximumSize | `
Format-Volume -FileSystem NTFS -NewFileSystemLabel "Quorum" -Confirm:$false
```

DriveLetter	FileSystemLabel	FileSystem	DriveType	HealthStatus	OperationalStatus	SizeRemaining	Size
H	Quorum	NTFS	Fixed	Healthy	OK	1.98 GB	2.01 GB

- On the same Hyper-V hosts, run the following PowerShell command to prepare the disk for Cluster Shared Volume (CSV) by initializing, creating new partition, and formatting it.

```
Get-Disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 2 | `
Initialize-Disk -PartitionStyle GPT -PassThru | `
New-Partition -AssignDriveLetter -UseMaximumSize | `
Format-Volume -FileSystem NTFS -NewFileSystemLabel "VMStore" -Confirm:$false
```

Figure 112 Initialize, Partition and Formatting Volumes

```
PS C:\Windows\system32> get-disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 2 | `
Initialize-Disk -PartitionStyle GPT -PassThru | `
New-Partition -AssignDriveLetter -UseMaximumSize | `
Format-Volume -FileSystem NTFS -NewFileSystemLabel "VMStore" -Confirm:$false
```

DriveLetter	FileSystemLabel	FileSystem	DriveType	HealthStatus	OperationalStatus	SizeRemaining	Size
I	VMStore	NTFS	Fixed	Healthy	OK	499.79 GB	499.94 GB

- Run the "Get-Disk" PowerShell cmdlet to verify as shown below .

Figure 113 Listing the Volumes

```
PS C:\Windows\system32> Get-Disk
```

Number	Friendly Name	Serial Number	HealthStatus	OperationalStatus	Total Size	Partition Style
0	NETAPP LUN C-Mode	808b5]Miw/xk	Healthy	Online	500.07 GB	GPT
2	NETAPP LUN C-Mode	808b5]Miw/xm	Healthy	Online	500.07 GB	GPT
1	NETAPP LUN C-Mode	808b5]Miw/xw	Healthy	Online	2.01 GB	MBR
4	Cisco CIMC-Mapped vHDD	20170330-4	Healthy	No Media	0 B	RAW
3	Cisco vKVM-Mapped vHDD	20170330-1	Healthy	No Media	0 B	RAW

- Run the following commands on the other Hyper-V hosts to prepare the disks for cluster use. Note that only the disks are brought online:

```
Get-Disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 1 | `
Set-Disk -IsOffline $false
```

Figure 114 Setting Quorum Disk Online

```
PS C:\Windows\system32> get-disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 1 | `
Set-disk -IsOffline $false
```

```
Get-Disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 2 | `
Set-Disk -IsOffline $false
```

Figure 115 Setting Other Disks Online

```
PS C:\Windows\system32> get-disk | where PartitionStyle -EQ 'RAW' | `
where Number -EQ 2 | `
Set-disk -IsOffline $false
```

## Test and Validate the Cluster

To test and validate the cluster, follow these steps:

1. Before you create the failover cluster, it is strongly recommended that you validate the configuration to make sure that the hardware and its settings are compatible with failover clustering. To run the cluster validation tests, run the following commands:

```
Test-Cluster -Node "FP-HV-Host1","FP-HV-Host2"
```

Figure 116 Test Windows Failover Cluster Nodes

```
PS C:\Windows\system32> Test-Cluster -Node "FP-HV-Host1","FP-HV-Host2"
```

Mode	LastWriteTime	Length	Name
-a----	12/18/2018 2:46 PM	867020	Validation Report 2018.12.18 At 14.43.07.htm

2. Open the failover cluster validation report and fix any errors/warnings before creating the failover cluster. The failover cluster validation report default location is "C:\Windows\Cluster\Reports\"

Figure 117 Windows Failover Cluster Validation Report



## Failover Cluster Validation Report

<b>Node:</b>	FP-HV-Host1.aflexpod.cisco.com	Validated
<b>Node:</b>	FP-HV-HOST2.aflexpod.cisco.com	Validated
<b>Started</b>	12/18/2018 2:43:07 PM	
<b>Completed</b>	12/18/2018 2:46:17 PM	

The Validate a Configuration Wizard must be run after any change is made to the configuration of the cluster or hardware. For more information, see <http://go.microsoft.com/fwlink/p/?LinkId=280145>.

### Results by Category

Name	Result Summary	Description
<a href="#">Hyper-V Configuration</a>		Success
<a href="#">Inventory</a>		Success
<a href="#">Network</a>		Success
<a href="#">Storage</a>		Success
<a href="#">System Configuration</a>		Success

### Create a New Failover Cluster

1. Create the Windows failover cluster using the following PowerShell command:

```
New-Cluster -Name FP-HVSQL -Node FP-HV-Host1,FP-HV-Host2 -StaticAddress 192.168.96.201
```

Figure 118 Windows Failover Cluster Creation

```
PS C:\Windows\system32> New-Cluster -Name FP-HVSQL -Node FP-HV-Host1,FP-HV-Host2 -StaticAddress 192.168.96.201

Name
----
FP-HVSQL
```

2. Verify the cluster status using the following PowerShell command or Failover Cluster Manager as shown below.

Figure 119 Windows Failover Cluster Verification

```


PS C:\Windows\system32> Get-ClusterGroup

Name                OwnerNode  State
-----
Available Storage  FP-HV-Host1 Online
Cluster Group      FP-HV-Host1 Online

PS C:\Windows\system32> Get-ClusterQuorum

Cluster              QuorumResource
-----
FP-HVSQL              Cluster Disk 1
    
```

Figure 120 Windows Failover Cluster Manager



### Post Cluster Creation Tasks

1. Find which disk is in the “Available Storage” OwnerGroup using the following PowerShell command to add to CSV. In this example, Cluster Disk 2 is in the Available Storage OwnerGroup.

Get-ClusterResource

Figure 121 Getting Cluster Resources

```
PS C:\Windows\system32> Get-ClusterResource
```

Name	State	OwnerGroup	ResourceType
Cluster Disk 1	Online	Cluster Group	Physical Disk
Cluster Disk 2	Online	Available Storage	Physical Disk
Cluster IP Address	Online	Cluster Group	IP Address
Cluster Name	Online	Cluster Group	Network Name
Virtual Machine Cluster WMI	Online	Cluster Group	Virtual Machine Cluster WMI

2. Add a disk in the Available Storage to CSV as shown below.

```
Add-ClusterSharedVolume -Name "Cluster Disk 2"
```

Figure 122 Adding CSV Volume

```
PS C:\Windows\system32> Add-ClusterSharedVolume -Name "Cluster Disk 2"
```

Name	State	Node
Cluster Disk 2	Online	FP-HV-HOST2

3. The CSV Block cache is enabled by default in Windows Server 2016 but the size of the Cache is set to zero as shown below.

Figure 123 Check CSV Cache

```
PS C:\Windows\system32> (Get-Cluster).BlockCacheSize
0
```

4. Assign an appropriate size to the CSV Block cache for better virtual machine performance. In the following example, the cmdlet reserves 512 MB of memory to the CSV cache on each node in the failover cluster. For more information about CSV cache, go to: <https://blogs.msdn.microsoft.com/clustering/2013/07/19/how-to-enable-csv-cache/>

Figure 124 Add CSV Cache Size

```
PS C:\Windows\system32> (Get-Cluster).BlockCacheSize = 512
PS C:\Windows\system32> (Get-Cluster).BlockCacheSize
512
```

5. Rename the cluster networks. This step is optional but helps in easy identification and troubleshooting.

```
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.96.0"}).Name = "Mgmt_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.97.0"}).Name = "LM_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.98.0"}).Name = "Cluster_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.13.0"}).Name = "iSCSI-A_Network"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.23.0"}).Name = "iSCSI-B_Network"
```

Figure 125 Rename Cluster Networks

```
PS C:\Windows\system32> (Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.96.0"}).Name = "Mgmt_Network"
PS C:\Windows\system32> (Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.97.0"}).Name = "LM_Network"
PS C:\Windows\system32> (Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.98.0"}).Name = "Cluster_Network"
PS C:\Windows\system32> (Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.13.0"}).Name = "iSCSI-A_Network"
PS C:\Windows\system32> (Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.23.0"}).Name = "iSCSI-B_Network"
```


6. Configure the cluster network roles as per their intended use and purpose using the following PowerShell command. Role 3 = Allow Client and Cluster traffic, Role 1 = Allow only Cluster traffic, and Role 0 = Do not allow any cluster traffic.

```
(Get-ClusterNetwork -Name "Mgmt_Network").Role = 3
(Get-ClusterNetwork -Name "Cluster_Network").Role = 1
(Get-ClusterNetwork -Name "LM_Network").Role = 1
(Get-ClusterNetwork -Name "iSCSI-A_Network").Role = 0
(Get-ClusterNetwork -Name "iSCSI-B_Network").Role = 0
```

Figure 126 Configure Network Roles

```
PS C:\Windows\system32> (Get-ClusterNetwork -Name "Mgmt_Network").Role = 3
PS C:\Windows\system32> (Get-ClusterNetwork -Name "Cluster_Network").Role = 1
PS C:\Windows\system32> (Get-ClusterNetwork -Name "LM_Network").Role = 1
PS C:\Windows\system32> (Get-ClusterNetwork -Name "iSCSI-A_Network").Role = 0
PS C:\Windows\system32> (Get-ClusterNetwork -Name "iSCSI-B_Network").Role = 0
```


Figure 127 Windows Failover Cluster Manager



7. Configure the Live Migration network as shown in below.

```
Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name MigrationExcludeNetworks -Value ([String]::Join(";", (Get-ClusterNetwork | Where-Object {$_.Name -ne "LM_Network"}).ID))
```

Figure 128 Configure the Live Migration network




### Create Clustered Virtual Machine for RHEL

To create a clustered virtual machine for RHEL, follow these steps:


1. Log into a Hyper-V cluster node and open the Failover Cluster Manager by navigating to Server Manager > Tools.
2. Right-click Roles, select Virtual Machines and click New Virtual Machine... as shown below.

Figure 129 Create Virtual Machine in Failover Cluster




3. Select the target cluster node for virtual machine creation.

**Figure 130 Select Hyper-V node for Virtual Machine**




4. In the Before You Begin section of New Virtual Machine wizard, click Next.
5. In the Specify Name and Location section of New Virtual Machine Wizard, enter a name for the virtual machine and CSV path in the location field and click Next.
6. In the Specify Generation section of the New Virtual Machine Wizard, choose the generation of the virtual machine and click Next.
7. In the Assign Memory section of the New Virtual Machine Wizard, specify the amount of memory to allocate for the VM.
8. In the Configure Networking section of the New Virtual Machine Wizard, click the drop-down list next to Connection and select the virtual switch.

**Figure 131 Choose Network Switch**




9. In the Connect Virtual Hard Disk section of New Virtual Machine Wizard, select Create a virtual hard disk and enter the Name and Location for the VHD as shown below. Location should be the CSV path.

Figure 132 Adding a Virtual Machine to Failover Cluster



10. In the Installation Options section of New Virtual Machine Wizard, select Install an operating System from a bootable RHEL image file.

Figure 133 Installation Options for New Virtual Machine




11. In the Summary section of New Virtual Machine Wizard, confirm and click Finish.

12. In the Summary section of High Availability Wizard, you can view the result and report. Click Finish.

13. When the virtual machine is created, open the Failover Cluster Manager, right-click the virtual machine and click Settings.


14. Select the Network Adapter added in step 8 and select the "Enable virtual LAN identification" check box enabled and enter the VLAN ID as shown below .

Figure 134 Configuring Virtual Machine Network Adapter for Management Traffic




- For in-guest iSCSI, add two more network adapters by Click on Add hardware, select Network Adapter and click Add. Select the virtual switch (SET in this case) and enter the appropriate VLAN IDs as shown below.

Figure 135 Configuration Virtual Machine Network Adapter for iSCSI Storage Traffic




- For Generation 2 virtual machines, the Secure Boot is enabled by default. For virtual machines running RHEL, select the Microsoft UEFI Certificate Authority option from the Template drop-down list as shown below . If secure boot is not required uncheck the Enable Secure Boot check box.

Figure 136 Enabling Secure Boot



17. Navigate to Server Manager > Tools > Failover Cluster Manager > Roles. Right-click the newly created virtual machine and click Connect and Start the virtual machine to begin the installation of RHEL.

Figure 137 Connect to Virtual Machine



## RHEL on Virtual Machine Installation and Configuration

---

This section explains how to install and configure RedHat Enterprise Linux version 7.4 on a virtual machine.

### Download Red Hat Enterprise Linux 7.4 image

To download the Red Hat Enterprise Linux 7.4, follow these steps:


1. Click the following link: [https://access.redhat.com/downloads/content/69/ver=/rhel---7/7.4/x86\\_64/product-software](https://access.redhat.com/downloads/content/69/ver=/rhel---7/7.4/x86_64/product-software).
2. You will need a login id and password from redhat.com to download the iso image.
3. Download the rhel-server-7.4-x86\_64-dvd.iso image

### Install Red Hat Enterprise Linux

To install RHEL 7.4 on the virtual machine, follow these steps.

1. Make sure the downloaded RHEL 7.4 iso image is attached to the virtual machine.
2. Open the console session of virtual machine. Use VMware vCenter for ESXi deployment -OR Hyper-V Manager for Hyper-V deployment to open the virtual machine console.
3. Power on the virtual machine to install the Operating System.
4. On the virtual machine console, using up arrow highlight Install Red Hat Enterprise Linux 7.4 and press enter.
5. On the language settings, choose appropriate language and click Continue.
6. On the Installation Summary screen, under Localization settings, click Date & Time and set the time zone by selecting the appropriate Region and City drop-down list. Click Done to go back to the previous screen.
7. Under System settings, click Installation Destination to select the disk for Operating System installation.
8. Under Local Standard Disks, select the boot disk. Choose Automatic partition option under Other Storage Option section. Click Done to go back to previous screen.
9. Under System settings, click Network and Hostname to set IP address and host name of the virtual machine.
10. Select the first network adapter and click Configure.
11. Provide appropriate IP address as shown below.

Figure 138 Configuring IP Address



12. Click Save to save the IP details.
13. Select the first network adapter and click Turn It On. Note that the remaining two network adapters, used for in-guest iscsi storage access, can be configured after OS is installed.
14. Enter Fully Qualified Domain Name (FQDN) of the guest vm and click Done to return to previous screen.
15. On the installation Summary screen, click Begin Installation to begin the Operating System installation.
16. On the User Settings screen, provide password for root user complete the OS installation.
17. After OS installation is complete, reboot the vm.
18. When vm is rebooted, verify if you can logon to RHEL vm using root user.
19. Make sure the appropriate hypervisor’s integration tools are installed and running on the RHEL virtual machines.
20. Subscribe the virtual machine to Red Hat Linux using your RHEL user id and password as shown below.

```
subscription-manager register --username <username> --password <password> --auto-attach
```

## RHEL OS Configuration for Higher Database Performance

This section provides the RHEL OS configuration guidelines to achieve better SQL Server database performance. For more details, go to: <https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-performance-best-practices?view=sql-server-2017>

1. Log into the RHEL virtual machine using root user to perform the following tasks.
2. Edit the `/etc/sysctl.conf` and append the following settings at the end of the file.

```
kernel.numa_balancing = 0
kernel.sched_min_granularity_ns = 10000000
kernel.sched_wakeup_granularity_ns = 15000000
vm.dirty_background_ratio = 10
vm.dirty_ratio = 40
vm.swappiness = 10
vm.max_map_count = 262144
```

3. Process (enable) the above changes by running the following command.

```
sysctl -p
```

4. Make sure the Transparent Huge Pages (THP) is enabled. It is enabled by default from RHEL 7 onwards. Run the following command to ensure THP is enabled.

```
cat /sys/kernel/mm/*transparent_hugepage/enable
```

If your output of above command looks like “[always] madvise never”, then it indicates that THP is enabled.

5. Install optional but recommended system monitoring packages as shown below. These packages are useful for monitoring system performance.

```
yum -y install dstat
yum -y install sysstat
```



Additional SQL Server tunings are detailed in section [SQL Server 2017 Installation](#).

---

## Join Linux Virtual Machine to Microsoft Windows Active Directory (AD) Domain

This section explains how to add a RHEL virtual machine to the Microsoft Windows Active Directory Domain.



**This step is optional. When a virtual machine is added to Windows Active Directory domain, users will be able to be authenticated and authorized by Active Directory services. This enables users from Active Directory domains to transparently access Linux systems and services without having to create Linux user ids on each virtual machine.**

---

To join a RHEL virtual machine to a Windows Active Directory Domain, follow these steps:

1. Log into the RHEL virtual machine using root user to perform the following tasks.
2. Make sure you have updated the DNS IP and domain details in the `/etc/resolv.conf` file.
3. Install required packages by running the following command.

```
yum install sssd oddjob oddjob-mkhomedirectory samba-common-tools
```

```
yum install realmd krb5-workstation
```

4. Join the RHEL virtual machine to the domain by running the following command. You must authenticate using an Active Directory account that has sufficient privileges in AD to join a new machine to the domain. When the virtual machine joins the AD, it creates a new computer account in AD, create the /etc/krb5.keytab host keytab file and configure the domain in /etc/sss/sss.conf file.


**Figure 139 Joining RHEL Virtual Machine to the Domain**

```
[root@SQLVM1 ~]# realm join aflexpod.cisco.com -U 'flexadmin@AFLEXPOD.CISCO.COM' -v
* Resolving: ldap_tcp.aflexpod.cisco.com
* Performing LDAP DSE lookup on: 10.1.166.9
* Successfully discovered: aflexpod.cisco.com
Password for flexadmin@AFLEXPOD.CISCO.COM:
* Required files: /usr/sbin/oddjobd, /usr/libexec/oddjob/mkhomedir, /usr/sbin/sss, /usr/bin/net
* LANG=C LOGNAME=root /usr/bin/net -s /var/cache/realmd/realmd-smb-conf.FOC3TZ -U flexadmin@AFLEXPOD.CISCO.COM
ads join aflexpod.cisco.com
Enter flexadmin@AFLEXPOD.CISCO.COM's password:DNS update failed: NT_STATUS_INVALID_PARAMETER

Using short domain name -- AFLEXPOD
Joined 'SQLVM1' to dns domain 'aflexpod.cisco.com'
No DNS domain configured for sqlvm1. Unable to perform DNS Update.
* LANG=C LOGNAME=root /usr/bin/net -s /var/cache/realmd/realmd-smb-conf.FOC3TZ -U flexadmin@AFLEXPOD.CISCO.COM
ads keytab create
Enter flexadmin@AFLEXPOD.CISCO.COM's password:
* /usr/bin/systemctl enable sssd.service
Created symlink from /etc/systemd/system/multi-user.target.wants/sss.service to /usr/lib/systemd/system/sss.service.
* /usr/bin/systemctl restart sssd.service
* /usr/bin/sh -c /usr/sbin/authconfig --update --enablesssd --enablesssdauth --enablemkhomedir --nostart && /usr/bin/systemctl enable oddjobd.service && /usr/bin/systemctl start oddjobd.service
* Successfully enrolled machine in realm
[root@SQLVM1 ~]#
```

5. When the vm joined the domain successfully, verify if you are able to connect to virtual machine using AD domain user with SSH as shown below.

**Figure 140 Connecting RHEL Virtual Machine with AD Domain User**



```
flexadmin@aflexpod.cisco.com@SQLVM01:~
login as: flexadmin@AFLEXPOD.CISCO.COM
flexadmin@AFLEXPOD.CISCO.COM@SQLVM01's password:
Creating home directory for flexadmin@AFLEXPOD.CISCO.COM.
[flexadmin@aflexpod.cisco.com@SQLVM01 ~]$ pwd
/home/flexadmin@aflexpod.cisco.com
[flexadmin@aflexpod.cisco.com@SQLVM01 ~]$
```

### Configure iSCSI Network Adapters for Storage Access

This section explains how to configure the iSCSI network adapters and In-Guest iSCSI initiator for direct NetApp storage access using iSCSI protocol.

Each SQL virtual machine is configured with two network adapters for accessing NetApp storage volumes using iSCSI protocol. To configuring the IP address on each iSCSI network adapter, follow these steps:

- 
-  To assign IP addresses to all the network adapters allocated to the virtual machine, gather the MAC addresses and VLAN IDs from virtual machine settings.
-

- Review the iSCSI network adapters shown below. ens224 and ens256 are used for iSCSI storage access. Network Adapter ens224 is used for accessing NetApp storage volumes using Fabric Interconnect A path and while ens256 uses Fabric Interconnect B path.

Figure 141 Listing Virtual Machine Network Adapters

```
[root@SQLVM01 network-scripts]# ip link list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT qlen 1
   link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: ens192: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT qlen 1000
   link/ether 00:50:56:aa:79:96 brd ff:ff:ff:ff:ff:ff
3: ens224: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT qlen 1000
   link/ether 00:50:56:aa:9d:32 brd ff:ff:ff:ff:ff:ff
4: ens256: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT qlen 1000
   link/ether 00:50:56:aa:07:f1 brd ff:ff:ff:ff:ff:ff
[root@SQLVM01 network-scripts]#
```

- Network adapter configuration files are in the path /etc/sysconfig/network-scripts/. Configure each iSCSI network adapter with appropriate IP address and subnet mask. It is recommended to set the Maximum Transfer Unit as 9000 for better storage throughput. The following figure shows the configuration for ens224 and ens256 adapters.

Figure 142 iSCSI Network Adapter Configuration

<pre>[root@SQLVM01 network-scripts]# pwd /etc/sysconfig/network-scripts [root@SQLVM01 network-scripts]# ls ifcfg-ens192  ifdown          ifdown-ipv6  ifdown-routes ifcfg-ens224  ifdown-bnep    ifdown-isdn  ifdown-sit ifcfg-ens256  ifdown-eth     ifdown-post  ifdown-Team ifcfg-lo      ifdown-ippool  ifdown-ppp   ifdown-TeamPort [root@SQLVM01 network-scripts]# cat ifcfg-ens224 TYPE=Ethernet PROXY_METHOD=none BROWSER_ONLY=no BOOTPROTO=none DEFROUTE=yes IPV4_FAILURE_FATAL=no NAME=ens224 UUID=f1c9f307-33ac-4745-9d3d-64fc5ef83e32 DEVICE=ens224 ONBOOT=yes IPADDR=192.168.12.52 PREFIX=24 MTU=9000 ZONE=public [root@SQLVM01 network-scripts]#</pre>	<pre>[root@SQLVM01 network-scripts]# pwd /etc/sysconfig/network-scripts [root@SQLVM01 network-scripts]# ls ifcfg-ens192  ifdown          ifdown-ipv6  ifdown-routes ifcfg-ens224  ifdown-bnep    ifdown-isdn  ifdown-sit ifcfg-ens256  ifdown-eth     ifdown-post  ifdown-Team ifcfg-lo      ifdown-ippool  ifdown-ppp   ifdown-TeamPort [root@SQLVM01 network-scripts]# cat ifcfg-ens256 TYPE=Ethernet PROXY_METHOD=none BROWSER_ONLY=no BOOTPROTO=none DEFROUTE=yes IPV4_FAILURE_FATAL=no NAME=ens256 UUID=5747265a-1426-4959-919f-cd47e4f33a15 DEVICE=ens256 ONBOOT=yes IPADDR=192.168.22.52 PREFIX=24 MTU=9000 ZONE=public [root@SQLVM01 network-scripts]#</pre>
--	--

- Disable ipv6 on all the interfaces by append following two lines to /etc/sysctl.conf file:

```
net.ipv6.conf.all.disable_ipv6 = 1
net.ipv6.conf.default.disable_ipv6 = 1
```

- To make above changes effective, run the following command:

```
sysctl -p
```

- Add the following line to the /etc/ssh/sshd\_config file to avoid SSH Xforwarding:

```
AddressFamily inet
```

- Restart the sshd for changes to get effect:

```
systemctl restart sshd
```

7. Verify the MTU setting by running the following command. If MTU is properly set at all layers (Fabric Interconnects, Nexus Switches and NetApp Storage), the virtual machine should be able to contact storage without fragmenting the packets as shown below.

Figure 143 Verifying MTU Setting Between Virtual Machines and Storage

```
[root@SQLVM01 /]# ping -M do -s 8972 192.168.12.18 -I ens224
PING 192.168.12.18 (192.168.12.18) from 192.168.12.52 ens224: 8972(9000) bytes of data.
8980 bytes from 192.168.12.18: icmp_seq=1 ttl=64 time=0.286 ms
8980 bytes from 192.168.12.18: icmp_seq=2 ttl=64 time=0.228 ms
^C
--- 192.168.12.18 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.228/0.257/0.286/0.029 ms
[root@SQLVM01 /]#
[root@SQLVM01 /]# ping -M do -s 8972 192.168.12.19 -I ens224
PING 192.168.12.19 (192.168.12.19) from 192.168.12.52 ens224: 8972(9000) bytes of data.
8980 bytes from 192.168.12.19: icmp_seq=1 ttl=64 time=0.275 ms
8980 bytes from 192.168.12.19: icmp_seq=2 ttl=64 time=0.232 ms
^C
--- 192.168.12.19 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.232/0.253/0.275/0.026 ms
[root@SQLVM01 /]#
[root@SQLVM01 /]# ping -M do -s 8972 192.168.22.19 -I ens256
PING 192.168.22.19 (192.168.22.19) from 192.168.22.52 ens256: 8972(9000) bytes of data.
8980 bytes from 192.168.22.19: icmp_seq=1 ttl=64 time=0.284 ms
8980 bytes from 192.168.22.19: icmp_seq=2 ttl=64 time=0.265 ms
^C
--- 192.168.22.19 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.265/0.274/0.284/0.019 ms
[root@SQLVM01 /]#
[root@SQLVM01 /]# ping -M do -s 8972 192.168.22.18 -I ens256
PING 192.168.22.18 (192.168.22.18) from 192.168.22.52 ens256: 8972(9000) bytes of data.
8980 bytes from 192.168.22.18: icmp_seq=1 ttl=64 time=0.267 ms
8980 bytes from 192.168.22.18: icmp_seq=2 ttl=64 time=0.252 ms
^C
--- 192.168.22.18 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.252/0.259/0.267/0.017 ms
[root@SQLVM01 /]# █
```

## Install and Configure NetApp Linux Unified Host Utilities, Software iSCSI Initiator and Multipath

This section explains how to install and configure the NetApp Linux Host utilities, software iSCSI initiator and multipathing.

### Install NetApp Linux Unified Host Utilities

NetApp Linux Unified Host Utilities software package includes the sanlun utility and the documentation. The sanlun utility helps to manage LUNs and HBAs. To install the NetApp host utilities, follow these steps:

1. For the interoperability matrix, go to: <http://mysupport.netapp.com/matrix> and download the appropriate NetApp Linux Unified Host utilities using your login credentials. For this setup, netapp\_linux\_unified\_host\_utilities-7-1.x86\_64.rpm was downloaded.
2. Upload the file to the virtual machine and install it using the following command. This automatically installs the sanlun utility.

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64.rpm
```

## Install and Configure Software iSCSI Initiator

This section explains how to install and configure the software iSCSI initiator which is used to connect the NetApp storage volumes directly bypassing VMWare Kernel networking stack. To install and configure iSCSI initiator, follow these steps:

1. Install iSCSI initiator packages by running the following command on the virtual machine.

```
yum install iscsi-initiator-utils -y
```

2. Change the initiator name as appropriate as shown below.

Figure 144 Change iSCSI Initiator Name

```
[root@SQLVM01 /]#
[root@SQLVM01 /]# cat /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.1994-05.com.redhat:sqlvm01.aflexpod.cisco.com
[root@SQLVM01 /]#
[root@SQLVM01 /]# █
```

3. Enable and start the iscsi initiator service by running following commands:

```
systemctl enable iscsi
```

```
systemctl enable iscsid
```

```
systemctl start iscsi
```

```
systemctl start iscsid
```

4. Edit /etc/iscsi/iscsid.conf file and change 'node.session.timeo.replacement\_timeout' parameter value from 120 to 5 as recommended NetApp Host Utilities installation documentation as shown below.

Figure 145 Configuring iSCSI Initiator

```
[root@SQLVM01 /]#
[root@SQLVM01 /]# cat /etc/iscsi/iscsid.conf | grep -i 'node.session.timeo.replacement_timeout'
node.session.timeo.replacement_timeout = 120
[root@SQLVM01 /]# vi /etc/iscsi/iscsid.conf
[root@SQLVM01 /]# cat /etc/iscsi/iscsid.conf | grep -i 'node.session.timeo.replacement_timeout'
node.session.timeo.replacement_timeout = 5
[root@SQLVM01 /]# █
```

5. Discover the iSCSI paths using iscsiadm utility as shown below.

Figure 146 Discovering iSCSI Paths

```
[root@SQLVM01 ~]# iscsiadm -m discovery -t st -p 192.168.12.18
192.168.12.18:3260,1035 iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5
192.168.22.19:3260,1039 iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5
192.168.12.19:3260,1038 iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5
192.168.22.18:3260,1036 iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5
[root@SQLVM01 ~]# █
```

- Verify that all storage SVMs IP addresses along with iqn are discovered as shown below. Write down the iqn name of the target SVM since it is required for the subsequent commands.

Figure 147 Verifying All Storage SVMs IP Address and iqn are Discovered

```
[root@SQLVM01 ~]#
[root@SQLVM01 ~]# iscsiadm -m discovery --print=1
SENDTARGETS:
DiscoveryAddress: 192.168.12.18,3260
Target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5
  Portal: 192.168.12.18:3260,1035
    Iface Name: default
  Portal: 192.168.22.19:3260,1039
    Iface Name: default
  Portal: 192.168.12.19:3260,1038
    Iface Name: default
  Portal: 192.168.22.18:3260,1036
    Iface Name: default
DiscoveryAddress: 192.168.12.19,3260
DiscoveryAddress: 192.168.13.18,3260
DiscoveryAddress: 192.168.13.19,3260
DiscoveryAddress: 192.168.23.18,3260
DiscoveryAddress: 192.168.22.18,3260
DiscoveryAddress: 192.168.22.19,3260
iSNS:
No targets found.
STATIC:
No targets found.
FIRMWARE:
No targets found.
[root@SQLVM01 ~]#
[root@SQLVM01 ~]#
```

- Log into all the targets that were discovered using the target iqn as show below.

Figure 148 Log into iSCSI Targets

```
[root@SQLVM01 ~]#
[root@SQLVM01 ~]# iscsiadm -m node -T iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5 -l
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.12.18,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.22.19,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.12.19,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.22.18,3260] (multiple)
Login to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.12.18,3260] successful.
Login to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.22.19,3260] successful.
Login to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.12.19,3260] successful.
Login to [iface: default, target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5, portal: 192.168.22.18,3260] successful.
[root@SQLVM01 ~]#
```


- Check if the iSCSI session has been established between the virtual machine and target SVM as shown below.

Figure 149 Verify iSCSI Session Between Virtual Machine and Target SVM

```
[root@SQLVM01 ~]# iscsiadm -m session -P3
iSCSI Transport Class version 2.0-870
version 6.2.0.874-10
Target: iqn.1992-08.com.netapp:sn.a78f2780f80d11e892c200a098a9fed2:vs.5 (non-flash)
Current Portal: 192.168.12.18:3260,1035
Persistent Portal: 192.168.12.18:3260,1035
*****
Interface:
*****
Iface Name: default
Iface Transport: tcp
Iface Initiatorname: iqn.1994-05.com.redhat:sqlvm01.aflexpod.cisco.com
Iface IPaddress: 192.168.12.52
Iface HWaddress: <empty>
Iface Netdev: <empty>
SID: 1
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
Internal iscsid Session State: NO CHANGE
*****
Timeouts:
*****
Recovery Timeout: 120
Target Reset Timeout: 30
LUN Reset Timeout: 30
Abort Timeout: 15
****
CHAP:
****
username: <empty>
password: *****
username_in: <empty>
password_in: *****
*****
```


9. Log into the NetApp OnCommand System Manager to create initiator group for the virtual machine. The iqn of virtual machine will be automatically appears in the initiator lists. Select the iqn of the sqlvm01 virtual machine and click Add Initiator as shown below.

Figure 150 Creating Initiator Group with Virtual Machine iqn



10. Create Volumes/LUNs and grant access to the initiator group created in the above step as shown below. Grant access to SQL data and Log LUNs which will be used for storing SQL Server database data and transaction log files.

Figure 151 Granting LUN Access to the Virtual Machine



11. Verify that the LUNs are visible inside the virtual machine by running lsblk command as shown below.

Figure 152 LUNs Visibility Inside the Virtual Machine

```
[root@SQLVM01 ~]# lsblk
NAME                MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
fd0                  2:0    1    4K  0 disk
sda                  8:0    0   40G  0 disk
├─sda1                8:1    0   200M  0 part /boot/efi
├─sda2                8:2    0    1G   0 part /boot
└─sda3                8:3    0  38.8G  0 part
   └─rhel_sqlvm1-root 253:0   0  34.8G  0 lvm  /
      └─rhel_sqlvm1-swap 253:1   0    4G   0 lvm  [SWAP]
sdb                  8:16   0  400G  0 disk
sdc                  8:32   0  400G  0 disk
sdd                  8:48   0  400G  0 disk
sde                  8:64   0  400G  0 disk
sdf                  8:80   0  150G  0 disk
sdg                  8:96   0  150G  0 disk
sdh                  8:112  0  150G  0 disk
sdi                  8:128  0  150G  0 disk
sr0                  11:0    1   3.8G  0 rom
[root@SQLVM01 ~]#
```

### Install and Configure Multipathing

To install and configure multipath inside the RHEL virtual machine, follow these steps:

1. Install required multipath packages by running the following commands.

```
yum install device-mapper -y
```

```
yum install device-mapper-multipath -y
```

2. If multipath.conf does not exist in /etc/ folder, copy the sample file from /usr/share/doc/device-mapper-multipath-0.4.9/. Run the following command to copy the file to /etc/ folder.

```
cp /usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf /etc/
```

3. Enable and Start the multipath service by running the following commands.

```
systemctl enable multipathd.service
```

```
systemctl start multipathd.service
```

4. Run the following command to list the device ids (WWIDs) which are required for editing multipath.conf file.

Figure 153 List Device IDs

```
[root@SQLVM01 /]# multipath -ll
3600a098038303862535d4d69772f7853 dm-3 NETAPP ,LUN C-Mode
size=150G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- 34:0:0:1 sdh 8:112 active ready running
|  |- 35:0:0:1 sdi 8:128 active ready running
|+- policy='service-time 0' prio=10 status=enabled
|  |- 33:0:0:1 sdf 8:80 active ready running
|  - 36:0:0:1 sdg 8:96 active ready running
3600a098038303862535d4d69772f784f dm-2 NETAPP ,LUN C-Mode
size=400G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- 34:0:0:0 sdd 8:48 active ready running
|  |- 35:0:0:0 sde 8:64 active ready running
|+- policy='service-time 0' prio=10 status=enabled
|  |- 33:0:0:0 sdb 8:16 active ready running
|  - 36:0:0:0 sdc 8:32 active ready running
```

5. Edit the /etc/multipath.conf file. Comment all the lines except the lines shown below. Update these lines with the corresponding wwids gathered in the last command 'multipath -ll'.

Figure 154 Updating Multipath.conf File

```
*multipath.conf
49 ## of devices, such as all scsi devices, you should use a devnode line.
50 ## However, if you want to blacklist specific devices, you should use
51 ## a wwid line. Since there is no guarantee that a specific device will
52 ## not change names on reboot (from /dev/sda to /dev/sdb for example)
53 ## devnode lines are not recommended for blacklisting specific devices.
54 ##
55 blacklist {
56     wwid 26353900f02796769
57     devnode "^ (ram|raw|loop|fd|md|dm-[sr|scd|st) [0-9] *"
58     devnode "^hd[a-z]"
59 }
60 multipaths {
61     multipath {
62         wwid 3600a098038303862535d4d69772f7853
63         alias sqllog
64         path_grouping_policy multibus
65         path_selector "round-robin 0"
66         failback manual
67         rr_weight priorities
68         no_path_retry 5
69         prio alua
70     }
71     multipath {
72         wwid 3600a098038303862535d4d69772f784f
73         alias sqldata
74         path_grouping_policy multibus
75         path_selector "round-robin 0"
76         failback manual
77         rr_weight priorities
78         no_path_retry 5
79         prio alua
80     }
81 }
82 #devices {
83 # device {
84 #     vendor "COMPAQ "
85 #     product "HSV110 (C)COMPAQ"
86 #     path_grouping_policy multibus
87 #     path_checker readsector0
88 #     path_selector "round-robin 0"
89 #     hardware_handler "0"
90 #     failback 15
91 #     rr_weight priorities
```

6. Restart the multipathd service by running the following command.

```
systemctl start multipathd.service
```

7. Run the multipath -ll command and verify the alias is visible as shown below.

Figure 155 Device Alias Name

```
[root@SQLVM01 ~]# multipath -ll
sqldata (3600a098038303862535d4d69772f784f) dm-2 NETAPP ,LUN C-Mode
size=400G features='4 queue_if_no_path pg_init_retries 50 retain_attached_hw_handle' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=30 status=active
|- 33:0:0:0 sdb 8:16 active ready running
|- 34:0:0:0 sdc 8:32 active ready running
|- 35:0:0:0 sde 8:64 active ready running
|- 36:0:0:0 sdd 8:48 active ready running
sqllog (3600a098038303862535d4d69772f7853) dm-3 NETAPP
size=150G features='4 queue_if_no_path pg_init_retries
+- policy='round-robin 0' prio=30 status=active
|- 34:0:0:1 sdg 8:96 active ready running
|- 33:0:0:1 sdf 8:80 active ready running
|- 36:0:0:1 sdi 8:128 active ready running
|- 35:0:0:1 sdh 8:112 active ready running
[root@SQLVM01 ~]#
```

```
[root@SQLVM01 /]#
[root@SQLVM01 /]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
fd0 2:0 1 4K 0 disk
sda 8:0 0 40G 0 disk
├─sda1 8:1 0 200M 0 part /boot/efi
├─sda2 8:2 0 1G 0 part /boot
├─sda3 8:3 0 38.8G 0 part
├─rhel_sqlvm1-root 253:0 0 34.8G 0 lvm /
├─rhel_sqlvm1-swap 253:1 0 4G 0 lvm [SWAP]
sdb 8:16 0 400G 0 disk
├─sqllog 253:2 0 400G 0 mpath
sdc 8:32 0 400G 0 disk
├─sqllog 253:2 0 400G 0 mpath
sdd 8:48 0 400G 0 disk
├─sqllog 253:2 0 400G 0 mpath
sde 8:64 0 400G 0 disk
├─sqllog 253:2 0 400G 0 mpath
sdf 8:80 0 150G 0 disk
├─sqldata 253:3 0 150G 0 mpath
sdg 8:96 0 150G 0 disk
├─sqldata 253:3 0 150G 0 mpath
sdh 8:112 0 150G 0 disk
├─sqldata 253:3 0 150G 0 mpath
sdi 8:128 0 150G 0 disk
├─sqldata 253:3 0 150G 0 mpath
sr0 11:0 1 3.8G 0 rom
[root@SQLVM01 /]#
```

8. The multipath SAN LUNs can be viewed using sanlun tool as shown below.

Figure 156 Viewing LUNs Using sanlun Utility

```
[root@SQLVM01 ~]#
[root@SQLVM01 ~]# sanlun lun show
controller(7mode/E-Series)/
vserver(cDOT/FlashRay)
-----
lun-pathname          device          host          lun          product
filename              adapter        protocol      size
-----
ESXi-Work-SVM        /vol/sql_esxi_log/SQL-ESXi-Log-01 /dev/sdh      host34       iSCSI       150g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_log/SQL-ESXi-Log-01 /dev/sdi      host35       iSCSI       150g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_log/SQL-ESXi-Log-01 /dev/sdf      host33       iSCSI       150g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_log/SQL-ESXi-Log-01 /dev/sdg      host36       iSCSI       150g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_data/SQL-ESXi-Data-01 /dev/sde      host35       iSCSI       400g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_data/SQL-ESXi-Data-01 /dev/sdd      host34       iSCSI       400g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_data/SQL-ESXi-Data-01 /dev/sdb      host33       iSCSI       400g       cDOT
ESXi-Work-SVM        /vol/sql_esxi_data/SQL-ESXi-Data-01 /dev/sdc      host36       iSCSI       400g       cDOT
[root@SQLVM01 ~]#
```

## Partition and Format iSCSI Volumes

To partition and format the discovered iSCSI devices, follow these steps:

1. Device-mapper-multipath service creates a mapper for each discovered device as shown below. In this example, the sqldata partition is created for 400G data LUN and sqllog partition is created for 150G log LUN.

Figure 157 DM-Multipath Mapper for Discovered iSCSI Devices

```
[root@SQLVM01 ~]# ls -lrt /dev/mapper/*
crw----- . 1 root root 10, 236 Dec 15 12:15 /dev/mapper/control
lrwxrwxrwx. 1 root root    7 Dec 15 12:15 /dev/mapper/rhel_sqlvm1-swap -> ../dm-1
lrwxrwxrwx. 1 root root    7 Dec 15 12:15 /dev/mapper/rhel_sqlvm1-root -> ../dm-0
lrwxrwxrwx. 1 root root    7 Dec 15 12:15 /dev/mapper/sqllog -> ../dm-3
lrwxrwxrwx. 1 root root    7 Dec 15 12:15 /dev/mapper/sqldata -> ../dm-2
[root@SQLVM01 ~]#
```

2. List and verify that all the multipath devices are visible by running following command.

```
dmsetup ls --target=multipath
```

Figure 158 List and verify Multipath Devices

```
[root@SQLVM01 ~]# dmsetup ls --target=multipath
sqlbackup (253, 4)
sqldata (253, 2)
sqllog (253, 3)
[root@SQLVM01 ~]#
```

3. Partition each device using fdisk utility. Run the following commands to partition each device.

```
fdisk /dev/mapper/sqldata
```

press n for creating new partition

press p to create primary partition

press 1 to create first primary partition on the device

press enter to start partition from 2048 sector

press p to view the newly created partition and then press w to save the changes.

4. Run the kpartx command against multipath device to create the partition mapping as shown below.

Figure 159 Partition Mapping Against Multipath Device using kpartx

```

[root@SQLVM01 ~]#
[root@SQLVM01 ~]# kpartx -a -v /dev/mapper/sqldata
add map sqldata1 (253:4): 0 838858752 linear /dev/mapper/sqldata 2048
[root@SQLVM01 ~]# █

[root@SQLVM01 ~]# lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
fd0                                  2:0    1    4K  0 disk
sda                                  8:0    0   40G  0 disk
├─sda1                               8:1    0   200M  0 part  /boot/efi
├─sda2                               8:2    0    1G  0 part  /boot
├─sda3                               8:3    0  38.8G  0 part
│   └─rhel_sqlvm1-root              253:0    0  34.8G  0 lvm    /
│       └─rhel_sqlvm1-swap          253:1    0    4G  0 lvm    [SWAP]
└─sdb                                8:16   0  400G  0 disk
    └─sqldata                       253:2    0  400G  0 mpath
        └─sqldata1                  253:4    0  400G  0 part
sdc                                  8:32   0  400G  0 disk
    └─sqldata                       253:2    0  400G  0 mpath
        └─sqldata1                  253:4    0  400G  0 part
sdd                                  8:48   0  400G  0 disk
    └─sqldata                       253:2    0  400G  0 mpath
        └─sqldata1                  253:4    0  400G  0 part
sde                                  8:64   0  400G  0 disk
    └─sqldata                       253:2    0  400G  0 mpath
        └─sqldata1                  253:4    0  400G  0 part
sdf                                  8:80   0  150G  0 disk
    └─sqllog                        253:3    0  150G  0 mpath
sdg                                  8:96   0  150G  0 disk
    └─sqllog                        253:3    0  150G  0 mpath
sdh                                  8:112  0  150G  0 disk
    └─sqllog                        253:3    0  150G  0 mpath
sdi                                  8:128  0  150G  0 disk
    └─sqllog                        253:3    0  150G  0 mpath
sr0                                  11:0    1   3.8G  0 rom
[root@SQLVM01 ~]# █

```

5. Repeat steps 1 to 4 to partition the sqllog multipath device.
6. Create XFS filesystem on the two partitions using the mkfs command as shown below.

Figure 160 Creating xfs Partitions

```
[root@SQLVM01 ~]#
[root@SQLVM01 ~]# mkfs.xfs /dev/mapper/sqldata1
meta-data=/dev/mapper/sqldata1  isize=512    agcount=4, agsize=26214336 blks
        =                               sectsz=4096  attr=2, projid32bit=1
        =                               crc=1      finobt=0, sparse=0
data      =                               bsize=4096 blocks=104857344, imaxpct=25
        =                               sunit=0    swidth=0 blks
naming    =version 2                       bsize=4096 ascii-ci=0 ftype=1
log       =internal log                   bsize=4096 blocks=51199, version=2
        =                               sectsz=4096 sunit=1 blks, lazy-count=1
realtime  =none                            extsz=4096 blocks=0, rtextents=0
[root@SQLVM01 ~]# mkfs.xfs /dev/mapper/sqllog1
meta-data=/dev/mapper/sqllog1  isize=512    agcount=4, agsize=9830336 blks
        =                               sectsz=4096  attr=2, projid32bit=1
        =                               crc=1      finobt=0, sparse=0
data      =                               bsize=4096 blocks=39321344, imaxpct=25
        =                               sunit=0    swidth=0 blks
naming    =version 2                       bsize=4096 ascii-ci=0 ftype=1
log       =internal log                   bsize=4096 blocks=19199, version=2
        =                               sectsz=4096 sunit=1 blks, lazy-count=1
realtime  =none                            extsz=4096 blocks=0, rtextents=0
[root@SQLVM01 ~]#
```

7. Create directories for mounting the devices that were created in the steps above. Run the following commands to create two directories, one for each partition.

```
mkdir /sqldbdata
```

```
mkdir /sqldblog
```

8. Find the block ids of the two partitions using the blkid command as shown below.

Figure 161 Find the Block Ids of the Two Partitions

```
[root@SQLVM01 ~]# blkid
/dev/sda1: SEC_TYPE="msdos" UUID="0E6E-0AAD" TYPE="vfat" PARTLABEL="EFI System Partition" PARTUUID="00b5c74c-9289-49f3-b8d2-205b0f58089b"
/dev/sda2: UUID="a8102954-a43c-4d79-ad2d-7128a2316024" TYPE="xfs" PARTUUID="7282529e-1132-4cc6-ac50-2fc37a668408"
/dev/sda3: UUID="Yit5AC-fWvd-I2Qp-3pFw-qtKc-x4MM-R2SwnG" TYPE="LVM2_member" PARTUUID="c0cf0104-2449-4bb2-8904-863b04bd0a08"
/dev/sr0: UUID="2017-07-11-01-39-24-00" LABEL="RHEL-7.4 Server.x86_64" TYPE="iso9660" PTTYPE="dos"
/dev/mapper/rhel_sqlvm1-root: UUID="a9b7854d-c40d-445b-a1a6-6a393f2fa2ed" TYPE="xfs"
/dev/mapper/rhel_sqlvm1-swap: UUID="705a9610-1c6a-4b1e-b4a2-b9475b8f8650" TYPE="swap"
/dev/mapper/sqldata: PTTYPE="dos"
/dev/mapper/sqllog: PTTYPE="dos"
/dev/mapper/sqldata1: UUID="a20ab86c-0eae-4d80-af55-7473d143da42" TYPE="xfs"
/dev/mapper/sqllog1: UUID="e609b3e2-8137-4221-b83e-e0115d40b0c1" TYPE="xfs"
[root@SQLVM01 ~]#
```

9. Edit the /etc/fstab file using the vi editor to add the mount entries and then verify the partitions using the cat command as shown below.

Figure 162 Adding Fstab Entries

```
[root@SQLVM01 ~]# cat /etc/fstab
#
# /etc/fstab
# Created by anaconda on Thu Dec 13 05:13:59 2018
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
/dev/mapper/rhel_sqlvm1-root / xfs defaults 0 0
UUID=a8102954-a43c-4d79-ad2d-7128a2316024 /boot xfs defaults 0 0
UUID=0E6E-0AAD /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/mapper/rhel_sqlvm1-swap swap swap defaults 0 0
#
UUID=a20ab86c-0eae-4d80-af55-7473d143da42 /sqldbdata xfs defaults,noatime,_netdev 0 0
UUID=e609b3e2-8137-4221-b83e-e0115d40b0c1 /sqldblog xfs defaults,noatime,_netdev 0 0
[root@SQLVM01 ~]#
```

10. Run the following command to mount the two devices to corresponding mount points as mentioned in the /etc/fstab file.

```
mount -a
```

## SQL Server 2017 on RHEL Installation and Configuration

This section explains how to install and configure SQL Server 2017 on RHEL virtual machine. For more information, go to: <https://docs.microsoft.com/en-us/sql/linux/quickstart-install-connect-red-hat?view=sql-server-2017>

### SQL Server 2017 Installation

To install SQL Server 2017 on the RHEL virtual machine, follow these steps:

1. Log into virtual machine as root user.
2. Install Microsoft SQL Server repo files required for installing SQL Server 2017.

```
curl -o /etc/yum.repos.d/mssql-server.repo
https://packages.microsoft.com/config/rhel/7/mssql-server-2017.repo
```

3. Execute the following command to install SQL Server 2017.

```
yum install -y mssql-server
```

4. After installation completes, run mssql-conf to select required edition and set the password for SA login. For this documentation, Evaluation Edition (1) is chosen.

```
/opt/mssql/bin/mssql-conf setup
```

5. When setup is complete, enable SQL Server service and verify that SQL service is running by running the following commands.

```
systemctl enable mssql-server
```

```
systemctl status mssql-server
```

6. Allow the remote connections to SQL Server instance by opening SQL Server ports on the firewall. Execute the following commands to open SQL Server default port 1433 permanently.

```
firewall-cmd --zone=public --add-port=1433/tcp --permanent
```

```
firewall-cmd -reload
```

7. Install SQL Server command-line tools by running the following commands.

```
sudo curl -o /etc/yum.repos.d/msprod.repo
https://packages.microsoft.com/config/rhel/7/prod.repo
```

```
yum install -y mssql-tools unixODBC-devel
```

8. Set the path to SQL Server sqlcmd tool so that you do not have to specify the full path when using sqlcmd tool.

```
echo 'export PATH="$PATH:/opt/mssql-tools/bin"' >> ~/.bash_profile
```

```
echo 'export PATH="$PATH:/opt/mssql-tools/bin"' >> ~/.bashrc
```

```
source ~/.bashrc
```

9. Verify that you are able to connect to SQL Server locally using sqlcmd tool.

```
sqlcmd -S localhost -U sa -P <sa login Password>
```

## Change Memory Settings

To make there is enough free physical memory for the Linux Operating System, the SQL Server process uses only 80 percent of the physical RAM by default. However, the default behavior may not be appropriate for all the systems. It is recommended to manually configure the memory for SQL Server using `mssql-conf` tool. When changing this setting, make sure there is enough memory for Linux OS and other applications running on the same virtual machine. Figure 163 shows setting 12G memory for the SQL Server on a virtual machine configured with 16G memory.

Figure 163 Setting Memory for SQL Server

```
[root@SQLVM01 ~]# /opt/mssql/bin/mssql-conf set memory.memorylimitmb 12288
SQL Server needs to be restarted in order to apply this setting. Please run
'systemctl restart mssql-server.service'.
[root@SQLVM01 ~]# systemctl restart mssql-server.service
[root@SQLVM01 ~]# cat /var/opt/mssql/mssql.conf
[sqlagent]
enabled = false

[EULA]
accepteula = Y

[network]
kerberoskeytabfile = /var/opt/mssql/secrets/mssql.keytab

[memory]
memorylimitmb = 12288

[hadr]
hadrenabled = 1

[root@SQLVM01 ~]#
```

## Change the Default Directories for Databases

By default, SQL Server stores the database files in the default location (`/var/opt/mssql/data/`) until unless location of the database files is mentioned. It is a best practice to change the default database files location to point to the multipath devices that were previously created. To change the default database files location to the newly mounted multipath device, follow these steps:

1. Change the owner and group of the directories (mountpoints) where the NetApp devices are mounted.
2. Change the default data and log directories as shown below.


Figure 164 Changing Default Data and Log Directories

```
[root@SQLVM01 /]#
[root@SQLVM01 /]# chown mssql /sqldbdata/
[root@SQLVM01 /]# chgrp mssql /sqldbdata/
[root@SQLVM01 /]# chown mssql /sqldblog/
[root@SQLVM01 /]# chgrp mssql /sqldblog/
[root@SQLVM01 /]#
[root@SQLVM01 /]# /opt/mssql/bin/mssql-conf set filelocation.defaultdatadir /sqldbdata/
SQL Server needs to be restarted in order to apply this setting. Please run
'systemctl restart mssql-server.service'.
[root@SQLVM01 /]# sudo /opt/mssql/bin/mssql-conf set filelocation.defaultlogdir /sqldblog/
SQL Server needs to be restarted in order to apply this setting. Please run
'systemctl restart mssql-server.service'.
[root@SQLVM01 /]# systemctl restart mssql-server
[root@SQLVM01 /]#
```

## Configure Tempdb Database Files

The SQL Server on Linux installation does not offer an option to configure multiple tempdb files during the installation, it is recommended to create multiple tempdb data files after the installation. Figure 165 shows eight configured data files and one log file on the multipath devices. The SQL Server Management studio can be used to configure the TempDB database.

Figure 165 Configure TempDB database files



```
SQLQuery8.sql - S...POD\sqldadmin (61)*  SQLQuery7.sql - S...POD\sqldadmin (61)*  SQLQuery6.sql - L...FLEXPOD\sqldadmin)*  SQLQuery4.sql - S...POD\sqldadmin (61)*
sp_helpdb tempdb
go

alter database tempdb modify file ( NAME=tempdev , filename='/sqldbdata/tempdb.mdf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb modify file ( NAME=templog , filename='/sqldblog/templog.ldf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp2 , filename='/sqldbdata/temp2.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp3 , filename='/sqldbdata/temp3.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp4 , filename='/sqldbdata/temp4.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp5 , filename='/sqldbdata/temp5.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp6 , filename='/sqldbdata/temp6.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp7 , filename='/sqldbdata/temp7.ndf', size=8192 KB ,FILEGROWTH=64MB)
alter database tempdb ADD FILE ( NAME=temp8 , filename='/sqldbdata/temp8.ndf', size=8192 KB ,FILEGROWTH=64MB)

100 %
Messages
The file "tempdev" has been modified in the system catalog. The new path will be used the next time the database is started.
The file "templog" has been modified in the system catalog. The new path will be used the next time the database is started.
```

## Add SQL Server to Active Directory Domain

Active Directory authentication enables domain-joined clients on either Windows or Linux to authenticate to SQL Server using their domain credentials and the Kerberos protocol.


AD Authentication has the following advantages over SQL Server Authentication:

- Users authenticate through single sign-on, without being prompted for a password.
- By creating logins for AD groups, you can manage access and permissions in SQL Server using AD group memberships.
- Each user has a single identity across your organization, so it eliminates tracking of SQL Server logins created against which user.
- AD enables you to enforce a centralized password policy across your organization.

For detailed information to configure the SQL Server on Linux for Active Directory Authentication, go to: <https://docs.microsoft.com/en-us/sql/linux/sql-server-linux-active-directory-authentication?view=sql-server-2017#createuser>

Below is an example of integrated authentication in Windows VM.

Figure 166 Connecting SQL Server with Integrated Authentication in Windows VM



```
flexadmin@aflexpod.cisco.com@SQLVM01:~  
flexadmin@aflexpod.cisco.com@SQLVM01 ~]$ whoami  
flexadmin@aflexpod.cisco.com  
flexadmin@aflexpod.cisco.com@SQLVM01 ~]$ sqlcmd -E -S localhost  
1> select @@version  
2> go  
  
-----  
-----  
Microsoft SQL Server 2017 (RTM-CU12) (KB4464082) - 14.0.3045.24 (X64)  
    oct 18 2018 23:11:05  
    Copyright (C) 2017 Microsoft Corporation  
    Enterprise Evaluation Edition (64-bit) on Linux (Red Hat Enterprise Linux)  
  
(1 rows affected)  
1> exit  
flexadmin@aflexpod.cisco.com@SQLVM01 ~]$ sqlcmd -E -S SQLVM02  
1> █
```

## SQL Server Always On Availability Group on RHEL Virtual Machines Deployment

This section details the process to deploy SQL Server Always On Availability Groups (AG) for high availability of databases on RHEL virtual machines using pacemaker cluster technology.

SQL Server Always On Availability Groups on Linux environments leverages underlying clustering technology called pacemaker. Pacemaker is a robust and powerful open source cluster resource manager which is shipping with Red Hat Enterprise Linux 7 as high availability add on. For more details on pacemaker cluster, refer: [https://access.redhat.com/documentation/en-us/red\\_hat\\_enterprise\\_linux/7/pdf/high\\_availability\\_add-on\\_reference/Red\\_Hat\\_Enterprise\\_Linux-7-High\\_Availability\\_Add-On\\_Reference-en-US.pdf](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/pdf/high_availability_add-on_reference/Red_Hat_Enterprise_Linux-7-High_Availability_Add-On_Reference-en-US.pdf)

For achieving high availability with automatic failover of SQL Server databases using Always On Availability Groups on RHEL requires minimum of three RHEL hosts. The following sections provides steps for deploying SQL Server Always On Availability Groups on three RHEL virtual machines.

### Install and Configure Pacemaker Cluster on RHEL Virtual Machines

To deploy and configure the pacemaker cluster, follow these steps:

1. Log into the RHEL hosts (SQLVM01, SQLVM02 and SQLVM03) using root user.
2. Enter the fully qualified names of all the cluster participating hosts in the /etc/hosts file on all the hosts as shown below for SQLVM01. This helps to resolve the IP addresses locally.



This step is not required if all the RHEL hosts are registered with their IP addresses in the DNS server.

Figure 167 Updating Fully Qualified Names

```
[root@SQLVM01 /]# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.94.51 SQLVM01.aflexpod.cisco.com    SQLVM01
192.168.94.52 SQLVM02.aflexpod.cisco.com    SQLVM02
192.168.94.53 SQLVM03.aflexpod.cisco.com    SQLVM03
[root@SQLVM01 /]#
```



The cluster participating hosts do not need to be part of Active Directory. Pacemaker also works with RHEL hosts which are not part of any domain

3. The pacemaker cluster comes as an add-on package with the RHEL 7.4 ISO installation media. To mount and configure the yum repository list using local media, run the following steps:



Before running the following commands, make sure RHEL 7.4 DVD is added to the virtual machines. Run the following step on all RHEL hosts.

```
mount -o loop /dev/sr0 /mnt
cp /mnt/media.repo /etc/yum.repos.d/rhel7dvd.repo
chmod 644 /etc/yum.repos.d/rhel7dvd.repo
```

4. Edit the `/etc/yum.repos.d/rhel7dvd.repo` file with the information shown below and save the file. Run the following step on all RHEL hosts.

Figure 168 Configuring rheldvd.repo File

```
[root@SQLVM01 /]#
[root@SQLVM01 /]# cat /etc/yum.repos.d/rhel7dvd.repo
[InstallMedia]
name=Red Hat Enterprise Linux 7.4
baseurl=file:///mnt/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
gpgcheck=0
enabled=1

[repo-ha]
gpgcheck=0
enabled=1
baseurl=file:///mnt/addons/HighAvailability
name=repo-ha

[repo-storage]
gpgcheck=0
enabled=1
baseurl=file:///mnt/addons/ResilientStorage
name=repo-storage

[root@SQLVM01 /]#
```

5. Install the required packages by running the following command. Run the following commands on all RHEL hosts.

```
yum install pacemaker pcs fence-agents-all resource-agents
```

6. Enable and start the pacemaker and pcsd services on all hosts. Run the following commands on all RHEL hosts.

```
systemctl start pcsd.service
systemctl enable pcsd.service
systemctl start pacemaker
systemctl enable pacemaker
```

7. Open the firewall ports for pacemaker cluster communication. Run the following commands on all RHEL hosts.

```
firewall-cmd --permanent --add-service=high-availability
firewall-cmd --add-service=high-availability
firewall-cmd -reload
```

8. Set the password for the default user that is created when installing Pacemaker and Corosync packages. Use the same password on all nodes.

```
passwd hacluster
```

9. Authenticate all the hosts from each host as shown below from the SQLVM01 host.

Figure 169 Authentication of Hosts

```
[root@SQLVM01 ~]# pcs cluster auth SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com SQLVM01.aflexpod.cisco.com
Username: hacluster
Password:
SQLVM03.aflexpod.cisco.com: Authorized
SQLVM02.aflexpod.cisco.com: Authorized
SQLVM01.aflexpod.cisco.com: Authorized
[root@SQLVM01 ~]#
```

10. Create the pacemaker cluster by running the following command on a single RHEL host:

```
pcs cluster setup --name sqlavg-clus SQLVM01.aflexpod.cisco.com
SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com
```

```
[root@SQLVM01 ~]#
[root@SQLVM01 ~]# pcs cluster setup --name sqlavg-clus SQLVM01.aflexpod.cisco.com SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com
Destroying cluster on nodes: SQLVM01.aflexpod.cisco.com, SQLVM02.aflexpod.cisco.com, SQLVM03.aflexpod.cisco.com...
SQLVM01.aflexpod.cisco.com: Stopping Cluster (pacemaker)...
SQLVM02.aflexpod.cisco.com: Stopping Cluster (pacemaker)...
SQLVM03.aflexpod.cisco.com: Stopping Cluster (pacemaker)...
SQLVM02.aflexpod.cisco.com: Successfully destroyed cluster
SQLVM01.aflexpod.cisco.com: Successfully destroyed cluster
SQLVM03.aflexpod.cisco.com: Successfully destroyed cluster

Sending 'pacemaker_remote authkey' to 'SQLVM01.aflexpod.cisco.com', 'SQLVM02.aflexpod.cisco.com', 'SQLVM03.aflexpod.cisco.com'
SQLVM01.aflexpod.cisco.com: successful distribution of the file 'pacemaker_remote authkey'
SQLVM02.aflexpod.cisco.com: successful distribution of the file 'pacemaker_remote authkey' SQLVM02.aflexpod.cisco.com: successful distribut
ion of the file 'pacemaker_remote authkey'
SQLVM03.aflexpod.cisco.com: successful distribution of the file 'pacemaker_remote authkey'

Sending cluster config files to the nodes...
SQLVM01.aflexpod.cisco.com: Succeeded
SQLVM02.aflexpod.cisco.com: Succeeded
SQLVM03.aflexpod.cisco.com: Succeeded

Synchronizing pcsd certificates on nodes SQLVM01.aflexpod.cisco.com, SQLVM02.aflexpod.cisco.com, SQLVM03.aflexpod.cisco.com...
SQLVM01.aflexpod.cisco.com: Success
SQLVM03.aflexpod.cisco.com: Success
SQLVM02.aflexpod.cisco.com: Success
Restarting pcsd on the nodes in order to reload the certificates...
SQLVM02.aflexpod.cisco.com: Success
SQLVM01.aflexpod.cisco.com: Success
SQLVM03.aflexpod.cisco.com: Success
[root@SQLVM01 ~]#
```

11. Start the pacemaker cluster services (pacemaker, pcs, corosyn, and so on.) and enable them on all hosts by running the following command on a single node.

```
pcs cluster start --all
pcs cluster enable -all
```

12. The following commands provide various cluster level information and configuration details.

```
pcs status
pcs status corosync
corosync-cmapctl | grep members
corosyn-cfgtool -s
```

13. Validate cluster configuration by running the following command. Cluster configuration errors, if any, are reported by this command.

```
crm_verify -L -V
```



You may notice some warnings related to the cluster fencing configuration. By default, the pacemaker cluster requires STONITH (Shoot The Other Node In The Head) or fencing device to be configured for the cluster setup. Fencing is the process of isolating and preventing a misbehaving node from affecting the availability of the cluster. Pacemaker supports variety of fencing devices and plugins. For more information about configuring the fencing devices on VMware environment, go to

<https://access.redhat.com/solutions/917813>. STONITH configuration is not supported in Hyper-V deployments.

14. Disable STONITH as shown below. Notice that after disabling STONITH, the warning messages disappear.

Figure 170 Disabling STONITH

```
[root@SQLVM02 ~]# crm_verify -L -V
error: unpack_resources: Resource start-up disabled since no STONITH resources have been defined
error: unpack_resources: Either configure some or disable STONITH with the stonith-enabled option
error: unpack_resources: NOTE: Clusters with shared data need STONITH to ensure data integrity
Errors found during check: config not valid
[root@SQLVM02 ~]#
[root@SQLVM02 ~]# pcs property set stonith-enabled=false
[root@SQLVM02 ~]#
[root@SQLVM02 ~]# pcs property show stonith-enabled
Cluster Properties:
stonith-enabled: false
[root@SQLVM02 ~]# crm_verify -L -V
[root@SQLVM02 ~]#
```

15. Set the cluster property cluster-recheck-interval to 2 minutes to reduce the time interval in which the cluster tries to attempt to restart a failed resource.

```
pcs property set cluster-recheck-interval=2min
```

16. Set the cluster property start-failure-is-fatal to true to avoid the cluster from restarting a failed resource on the same cluster node.

```
pcs property set start-failure-is-fatal=true
```

17. Pacemaker cluster can be monitored live using the following command.

```
crm_mon
```

Figure 171 Monitoring Cluster

```
Stack: corosync
Current DC: SQLVM03.aflexpod.cisco.com (version 1.1.16-12.e17-94ff4df) - partition with quorum
Last updated: Mon Dec 17 01:03:56 2018
Last change: Mon Dec 17 00:57:58 2018 by root via cibadmin on SQLVM02.aflexpod.cisco.com


3 nodes configured
0 resources configured

Online: [ SQLVM01.aflexpod.cisco.com SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]

No active resources
```


18. Pacemaker cluster can also be created and managed by the User Interface (UI). To create or manage an existing cluster, enter the user IP address of any RHEL host along with port 2224 as shown below. Use the hacluster IP URL to log into the cluster. In this case, <http://192.168.94.53:2222/login>.

Figure 172 Create or Manage Pacemaker Cluster Using UI



The figure below shows the details of an existing cluster.

Figure 173 Managing Cluster Using Pacemaker UI



## Install and Configure Always On Availability Group on Pacemaker Cluster

To install and configure Always On Availability Groups (AG) and how to add and manage AG resource in the pacemaker cluster, follow these steps:

1. Enable SQL Server Always On Availability Group and restart SQL service on all RHEL hosts by running the following command.

```
/opt/mssql/bin/mssql-conf set hadr.hadrenabled 1
systemctl restart mssql-service
```

2. Install Linux cluster resource agent for SQL Server Always On Availability Group by running the following command on all the rhel hosts. After the installation complete, you will notice resource provider for AG is listed

```
yum install mssql-server-ha -y
pcs resource providers
```

3. Enable AlwaysOn\_health event session of SQL Server instance on all the rhel hosts. Use SQL Management studio and run the following command on each SQL instance.

```
ALTER EVENT SESSION AlwaysOn_health ON SERVER WITH (STARTUP_STATE=ON)
```

4. SQL Server service on Linux uses certificates to authenticate communication between the mirroring endpoints. The following Transact-SQL script creates a master key and a certificate. It then backs up the certificate and secures the file with a private key. Update the script with strong passwords. Connect to the primary SQL Server instance (in this case, SQLVM01) and run the following script to create the certificate.

```
CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'YourMasterKeyPwd!';
CREATE CERTIFICATE dbm_certificate WITH SUBJECT = 'dbm';
BACKUP CERTIFICATE dbm_certificate
    TO FILE = '/var/opt/mssql/data/dbm_certificate.cer'
    WITH PRIVATE KEY (
        FILE = '/var/opt/mssql/data/dbm_certificate.pvk',
        ENCRYPTION BY PASSWORD = 'YourPrivateKeyPwd!'
    );
```

5. The script (above) creates two files; dbm\_certificate.cer and dbm\_certificate.pvk in /var/opt/mssql/data/ location with mssql as their file owner as shown below.

**Figure 174 Creating Certificates On Primary Replica**

```
[root@SQLVM01 ~]# ll /var/opt/mssql/data/dbm*
-rw-rw----. 1 mssql mssql 667 Dec 17 06:33 /var/opt/mssql/data/dbm_certificate.cer
-rw-rw----. 1 mssql mssql 1212 Dec 17 06:33 /var/opt/mssql/data/dbm_certificate.pvk
[root@SQLVM01 ~]# █
```

6. Copy these two files to each secondary replica (SQLVM02 and SQLVM03) to the same location as shown below.

Figure 175 Copy Certificates to Secondary Replicas

```
[root@SQLVM01 ~]# cd //var/opt/mssql/data/
[root@SQLVM01 data]# scp dbm_certificate.* root@SQLVM02:/var/opt/mssql/data/
The authenticity of host 'sqlvm02 (192.168.94.52)' can't be established.
ECDSA key fingerprint is SHA256:oUTLZ+U5Fbr73vRS2rEmQ4sydgPM9cXnF9I4RYmK5Nc.
ECDSA key fingerprint is MD5:c2:f0:a4:e5:49:7e:b7:4d:cf:f3:ee:85:4e:40:d2:77.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'sqlvm02,192.168.94.52' (ECDSA) to the list of known hosts.
root@sqlvm02's password:
dbm_certificate.cer                                100% 667   492.0KB/s   00:00
dbm_certificate.pvk                                100% 1212  972.7KB/s   00:00
[root@SQLVM01 data]# scp dbm_certificate.* root@SQLVM03:/var/opt/mssql/data/
The authenticity of host 'sqlvm03 (192.168.94.53)' can't be established.
ECDSA key fingerprint is SHA256:oUTLZ+U5Fbr73vRS2rEmQ4sydgPM9cXnF9I4RYmK5Nc.
ECDSA key fingerprint is MD5:c2:f0:a4:e5:49:7e:b7:4d:cf:f3:ee:85:4e:40:d2:77.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'sqlvm03,192.168.94.53' (ECDSA) to the list of known hosts.
root@sqlvm03's password:
dbm_certificate.cer                                100% 667   398.3KB/s   00:00
dbm_certificate.pvk                                100% 1212  986.6KB/s   00:00
[root@SQLVM01 data]#
```

- These files are copied from primary to secondaries using the root user; the current owner of these files is root. Change the file owner of these files to mssql on each secondary replica as shown below, for example; Changing owner on SQLVM02.

Figure 176 Changing Owner of the Files

```
[root@SQLVM02 data]#
[root@SQLVM02 data]# ll dbm_certificate.*
-rw-r-----. 1 root root 667 Dec 17 06:40 dbm_certificate.cer
-rw-r-----. 1 root root 1212 Dec 17 06:40 dbm_certificate.pvk
[root@SQLVM02 data]# chown mssql:mssql /var/opt/mssql/data/dbm*
[root@SQLVM02 data]# ll dbm_certificate.*
-rw-r-----. 1 mssql mssql 667 Dec 17 06:40 dbm_certificate.cer
-rw-r-----. 1 mssql mssql 1212 Dec 17 06:40 dbm_certificate.pvk
[root@SQLVM02 data]#
```

- Run the following script to create certificates on each secondary replica (SQLVM02 and SQLVM03) using the files copied in step 7.

```
CREATE MASTER KEY ENCRYPTION BY PASSWORD = 'YourMasterKeyPwd!';
go
CREATE CERTIFICATE dbm_certificate
    FROM FILE = '/var/opt/mssql/data/dbm_certificate.cer'
    WITH PRIVATE KEY (FILE='/var/opt/mssql/data/dbm_certificate.pvk',
        DECRYPTION BY PASSWORD='YourPrivateKeyPwd!');
go
```

- Create the database mirroring endpoints all the replicas by running the following script.

```
CREATE ENDPOINT [Hadr_endpoint]
    AS TCP (LISTENER_PORT = 5022)
    FOR DATABASE_MIRRORING (
        ROLE = ALL,
```

```

AUTHENTICATION = CERTIFICATE dbm_certificate,
ENCRYPTION = REQUIRED ALGORITHM AES
);

```

```
ALTER ENDPOINT [Hadr_endpoint] STATE = STARTED;
```

10. Open the endpoint tcp ports on the firewall by running the following script on each replica.

```
firewall-cmd --zone=public --add-port=5022/tcp --permanent
```

```
firewall-cmd -reload
```

11. Create the Always On Availability Group by running the following script on the primary replica as shown below. Note that Cluster\_Type is set to External and seeding (initial synchronization) is set to automatic. The three RHEL hosts are configured with synchronous replication with automatic failover.

Figure 177 Creating Always On Availability Group

```

CREATE AVAILABILITY GROUP [LINUX_SQLAG1]
WITH (DB_FAILOVER = ON, CLUSTER_TYPE = EXTERNAL)
FOR REPLICA ON
N'SQLVM01'
WITH (
ENDPOINT_URL = N'TCP://SQLVM01.AFLEXPOD.CISCO.COM:5022',
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
FAILOVER_MODE = EXTERNAL,
SEEDING_MODE = AUTOMATIC
),
N'SQLVM02'
WITH (
ENDPOINT_URL = N'TCP://SQLVM02.AFLEXPOD.CISCO.COM:5022',
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
FAILOVER_MODE = EXTERNAL,
SEEDING_MODE = AUTOMATIC
),
N'SQLVM03'
WITH(
ENDPOINT_URL = N'TCP://SQLVM03.AFLEXPOD.CISCO.COM:5022',
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
FAILOVER_MODE = EXTERNAL,
SEEDING_MODE = AUTOMATIC
);

ALTER AVAILABILITY GROUP [LINUX_SQLAG1] GRANT CREATE ANY DATABASE;

```

100 %

Messages

Commands completed successfully.

100 %

Query executed successfully. | SQLVM01 (14.0 RTM) | sa (60)

12. To have automatic failover of SQL resources in the pacemaker cluster, mssql-server-ha resource agent needs a login to connect the SQL instances. Create a SQL login with the appropriate permissions by running this script on each replica.

```
USE [master]
```

```
GO
```

```
CREATE LOGIN [pacemakerLogin] with PASSWORD= N'password'
```

```
ALTER SERVER ROLE [sysadmin] ADD MEMBER [pacemakerLogin]
```

13. Save the pacemakerLogin credentials in /var/opt/mssql/secrets/ path on each replica as shown below for SQLVM01.

Figure 178 Saving PacemakerLogin Credentials On All Replicas

```
[root@SQLVM01 /]# echo 'pacemakerLogin' >> /pacemaker-passwd
[root@SQLVM01 /]# echo 'Nbv12345!' >> /pacemaker-passwd
[root@SQLVM01 /]# mv /pacemaker-passwd /var/opt/mssql/secrets/passwd
[root@SQLVM01 /]# chmod 400 /var/opt/mssql/secrets/passwd
[root@SQLVM01 /]# cat /var/opt/mssql/secrets/passwd
pacemakerLogin
Nbv12345!
[root@SQLVM01 /]# █
```

14. Join the secondary replicas to the Always On Availability Group by running the following script on them (SQLVM02 and SQLVM03).

```
ALTER AVAILABILITY GROUP [LINUX_SQLAG1] JOIN WITH (CLUSTER_TYPE = EXTERNAL);
ALTER AVAILABILITY GROUP [LINUX_SQLAG1] GRANT CREATE ANY DATABASE;
```

15. Add a listener IP address to the Availability Group by running the following T-SQLscript on the primary replica.

```
ALTER AVAILABILITY GROUP LINUX_SQLAG1
ADD LISTENER N'SQLAG1_LISTNER'
(WITH IP
((N'192.168.94.55', N'255.255.255.0')), PORT=1433
);
GO
```

16. Grant the appropriate permission to the pacemakerLogin on Availability Group in all replicas by running the following script.

```
GRANT ALTER, CONTROL, VIEW DEFINITION ON AVAILABILITY GROUP::LINUX_SQLAG1 TO
[pacemakerLogin]
GRANT VIEW SERVER STATE TO [pacemakerLogin]
```

17. Create a sample database with full recovery mode and back it up to meet the database requirements of the Availability Group. Execute the following T-SQL script on the primary replica. Change the backup location as appropriate.

Figure 179 Creating Sample Database for AG

```

SQLQuery15.sql - SQ...01.master (sa (70))*  SQLQuery14.sql - L...SQLVMs.master (sa)*  SQLQuery12.sql - SQ...01.master (sa (60))*  SQ
CREATE DATABASE [SQLAG1_DB1]
ON PRIMARY
( NAME = N'SQLAG1_DB1', FILENAME = N'/sqldbdata/SQLAG1_DB1.mdf' , SIZE = 8192KB , FILEGROWTH = 65536KB )
LOG ON
( NAME = N'SQLAG1_DB1_log', FILENAME = N'/sqldblog/SQLAG1_DB1_log.ldf' , SIZE = 8192KB , FILEGROWTH = 65536KB )
GO

ALTER DATABASE [SQLAG1_DB1] SET RECOVERY FULL
GO

BACKUP DATABASE [SQLAG1_DB1]
TO DISK = N'/sqldbbackup/SQLAG1_DB1.bak';

```

100 %

Messages

Processed 304 pages for database 'SQLAG1\_DB1', file 'SQLAG1\_DB1' on file 1.  
 Processed 3 pages for database 'SQLAG1\_DB1', file 'SQLAG1\_DB1\_log' on file 1.  
 BACKUP DATABASE successfully processed 307 pages in 0.073 seconds (32.855 MB/sec).

18. Add the sample database created in the step 17 to the Always On Availability Group by running the following script on the Primary replica. Since the seeding mode is automatic, the secondary replica will be synchronized automatically by the Availability Group itself. For larger databases, it is recommended to manually synchronize the databases on the all the replicas and then join the database to the Availability Group.

```
ALTER AVAILABILITY GROUP LINUX_SQLAG1 ADD DATABASE SQLAG1_DB1;
```

19. You can view the final configuration by viewing the Availability Group dashboard or by running the T-SQL script as shown below. Note the Listener is currently offline since the listener and Availability Group is not yet added to the pacemaker cluster as the cluster resource.

```

SELECT ag.name AS 'AG Name', ar.replica_server_name AS 'Replica Instance',
DB_NAME(dr_state.database_id) AS 'Database', Location = CASE

WHEN ar_state.is_local = 1 THEN N'LOCAL'

ELSE 'REMOTE' END,

Role = CASE

WHEN ar_state.role_desc IS NULL THEN N'DISCONNECTED'

ELSE ar_state.role_desc END, ar_state.connected_state_desc AS 'Connection State',
ar.availability_mode_desc AS 'Mode', dr_state.synchronization_state_desc AS 'State'
FROM ((sys.availability_groups AS ag JOIN sys.availability_replicas AS ar ON
ag.group_id = ar.group_id ) JOIN sys.dm_hadr_availability_replica_states AS ar_state
ON ar.replica_id = ar_state.replica_id) JOIN sys.dm_hadr_database_replica_states
dr_state ON ag.group_id = dr_state.group_id and dr_state.replica_id =
ar_state.replica_id;

GO

SELECT b.dns_name, a.ip_address, a.ip_subnet_mask, a.state_desc, b.port FROM
sys.availability_group_listener_ip_addresses a INNER JOIN
sys.availability_group_listeners b ON a.listener_id=b.listener_id

```

Figure 180 Checking Status of Availability Group

```

SELECT ag.name AS 'AG Name', ar.replica_server_name AS 'Replica Instance',
       dr_name(dr_state.database_id) AS 'Database',
       location = CASE WHEN ar_state.is_local = 1 THEN N'LOCAL' ELSE 'REMOTE' END,
       role = CASE WHEN ar_state.role_desc IS NULL THEN N'DISCONNECTED' ELSE ar_state.role_desc END,
       ar_state.connected_state_desc AS 'Connection State', ar.availability_mode_desc AS 'Mode',
       ar_state.synchronization_state_desc AS 'State'
FROM ((sys.availability_groups AS ag JOIN sys.availability_replicas AS ar ON ag.group_id = ar.group_id )
JOIN sys.dm_hadr_availability_replica_states AS ar_state ON ar.replica_id = ar_state.replica_id)
JOIN sys.dm_hadr_database_replica_states dr_state ON ag.group_id = dr_state.group_id and dr_state.replica_id = ar_state.replica_id;

SELECT b.dns_name, a.ip_address, a.ip_subnet_mask, a.state_desc, b.port FROM sys.availability_group_listener_ip_addresses a
INNER JOIN sys.availability_group_listeners b
ON a.listener_id=b.listener_id

```

AG Name	Replica Instance	Database	Location	Role	Connection State	Mode	State
LINUX_SQLAG1	SQLVM01	SQLAG1_DB1	LOCAL	PRIMARY	CONNECTED	SYNCHRONOUS_COMMIT	SYNCHRONIZED
LINUX_SQLAG1	SQLVM02	SQLAG1_DB1	REMOTE	SECONDARY	CONNECTED	SYNCHRONOUS_COMMIT	SYNCHRONIZED
LINUX_SQLAG1	SQLVM03	SQLAG1_DB1	REMOTE	SECONDARY	CONNECTED	SYNCHRONOUS_COMMIT	NOT SYNCHRONIZING

dns_name	ip_address	ip_subnet_mask	state_desc	port
SQLAG1_LISTNER	192.168.94.55	255.255.255.0	OFFLINE	1433

20. Create the pacemaker cluster resource for Availability Group Linux\_SQLAG1 as resource type master/slave.

```
pcs resource create LINUX_SQLAG ocf:mssql:ag ag_name=LINUX_SQLAG1 meta failure-timeout=30s master notify=true
```

21. Run the following command to create the virtual IP address resource for the Always On Availability Group listener.

```
pcs resource create SQLAG1_LISTNER_VIP ocf:heartbeat:IPaddr2 ip=192.168.94.55 cidr_netmask=24
```

22. Create a DNS entry for the virtual IP address created in the step (above) so that you can use the listener name to connect to Availability Group databases.

Figure 181 DNS entry for Availability Group Listener

Name	Type	Data	Timestamp
(same as parent folder)	Name Server (NS)	aflexpod-ad.aflexpod.cisco.com.	static
(same as parent folder)	Start of Authority (SOA)	[1281], aflexpod-ad.aflexpod.cisco.com., hostmaster.aflexpod.cisco.com.	static
FP-HV-HOST2	Host (A)	192.168.96.12	12/17/2018 4:00:00 AM
FP-HV-Host1	Host (A)	192.168.96.11	12/14/2018 6:00:00 AM
SQL-Win-01	Host (A)	192.168.94.61	12/14/2018 12:00:00 AM
sqllinux-client	Host (A)	192.168.94.60	12/14/2018 1:00:00 AM
SQLAG1_LISTNER	Host (A)	192.168.94.55	static
sqlvm03	Host (A)	192.168.94.53	
sqlvm02	Host (A)	192.168.94.52	
sqlvm01	Host (A)	192.168.94.51	
sql-vm-esx-03	Host (A)	192.168.94.13	
sql-vm-esx-02	Host (A)	192.168.94.12	
sql-vm-esx-01	Host (A)	192.168.94.11	
bb09-ts	Host (A)	192.168.166.6	
bb09-affa300	Host (A)	192.168.166.20	
bb09-affa300-2	Host (A)	192.168.166.19	
bb09-affa300-1	Host (A)	192.168.166.18	
bb09-affa300-2-sp	Host (A)	192.168.166.17	
bb09-affa300-1-sp	Host (A)	192.168.166.16	
aflexpod-http	Host (A)	192.168.166.150	
bb09-6332	Host (A)	192.168.166.15	
bb09-6332-b	Host (A)	192.168.166.14	
bb09-6332-a	Host (A)	192.168.166.13	
bb09-93180-b	Host (A)	192.168.166.12	
bb09-93180-a	Host (A)	192.168.166.11	
(same as parent folder)	Host (A)	10.1.166.9	
aflexpod-ad	Host (A)	10.1.166.9	
docker-mgmt	Host (A)	10.1.166.20	
afzal-jump	Host (A)	10.1.166.2	
linuxjh	Host (A)	10.1.166.120	
bb09-infra-svm	Host (A)	10.1.166.12	
IOM-CTRL	Host (A)	10.1.166.11	
infra-esxi-02	Host (A)	10.1.166.102	
infra-esxi-01	Host (A)	10.1.166.101	

**SQLAG1\_LISTNER Properties**

Host (A) | Security

Host (uses parent domain if left blank):

Fully qualified domain name (FQDN):

IP address:

Update associated pointer (PTR) record

OK Cancel Apply

23. You can use the name SQLAG1\_LISTNER to connect to the Availability Group databases as shown below.

Figure 182 Connecting to Availability Group Databases using Listener Name

The screenshot shows the SQL Server Enterprise Manager interface. On the left, the Object Explorer tree is expanded to show the 'SQLAG1\_LISTNER' server. Under 'Availability Groups', the 'LINUX\_SQLAG1 (Primary)' group is visible, containing three replicas: 'SQLVM01 (Primary)', 'SQLVM02 (Secondary)', and 'SQLVM03 (Secondary)'. The 'SQLAG1\_LISTNER' listener is also visible under 'Availability Group Listeners'. The main window shows a query window with the following SQL code:

```
select @@servername
go
```

The Results pane shows a single row with the value 'SQLVM01'. A status bar at the bottom indicates 'Query executed successfully.' and 'SQLAG1\_LISTNER (14.0 RTM)'.

24. Add the co-location constraint for Availability Group and virtual IP resource to make sure that they always co-locate on the same host.

```
pcs constraint colocation add SQLAG1_LISTNER_VIP LINUX_SQLAG-master INFINITY with-rsc-role=Master
```

25. The Order constraint needs to be added so that the cluster will always start the Availability Group resource first and then the virtual IP resource.

```
pcs constraint order promote LINUX_SQLAG-master then start SQLAG1_LISTNER_VIP
```

26. The complete cluster status can be viewed using the pcs status command as shown below.

```
pcs status
```

Figure 183 Complete Cluster Status

```
[root@SQLVM01 /]# pcs status
Cluster name: sqlavg-clus
Stack: corosync
Current DC: SQLVM03.aflexpod.cisco.com (version 1.1.16-12.e17-94ff4df) - partition with quorum
Last updated: Sun Jan  6 12:04:49 2019
Last change: Sun Jan  6 11:24:00 2019 by root via crm_resource on SQLVM01.aflexpod.cisco.com

3 nodes configured
4 resources configured

Online: [ SQLVM01.aflexpod.cisco.com SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]

Full list of resources:

Master/Slave Set: LINUX_SQLAG-master [LINUX_SQLAG]
Masters: [ SQLVM01.aflexpod.cisco.com ]
Slaves: [ SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]
SQLAG1_LISTNER_VIP (ocf::heartbeat:IPaddr2): Started SQLVM01.aflexpod.cisco.com

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
[root@SQLVM01 /]# pcs constraint show
Location Constraints:
Resource: LINUX_SQLAG-master
Enabled on: SQLVM01.aflexpod.cisco.com (score:INFINITY) (role: Started)
Ordering Constraints:
promote LINUX_SQLAG-master then start SQLAG1_LISTNER_VIP (kind:Mandatory)
Colocation Constraints:
SQLAG1_LISTNER_VIP with LINUX_SQLAG-master (score:INFINITY) (with-rsc-role:Master)
Ticket Constraints:
[root@SQLVM01 /]# █
```

## Automatic Failover Testing of Availability Group

This section details the process to test the automatic failover of SQL Server Always On Availability Group configured in the previous step on a three node RHEL cluster.

When the RHEL host hosting the primary AG replica is not available for any reason, the pacemaker cluster will initiate the failover of Availability Group resources from the failed node to another node automatically.

Many failure scenarios were tested on the primary replica and verified that the Availability Group resource can failover from a failed primary replica to other nodes. For example, powering off the primary replica (as shown below 'shutdown now' executed on SQLVM01) resulted in an automatic failover of resources to other survival nodes as shown below. During the automatic failover of resources from failed node to other survival nodes, the listener will not be available for a few seconds. This means the applications need to resubmit the failed queries.

Figure 184 Initiating Failure on Primary Replica

```
[root@SQLVM01 ~]# pcs status
Cluster name: sqlavg-clus
Stack: corosync
Current DC: SQLVM01.aflexpod.cisco.com (version 1.1.16-12.e17-94ff4df) - partition with quorum
Last updated: Mon Jan 7 01:20:19 2019
Last change: Mon Jan 7 00:56:35 2019 by root via crm_resource on SQLVM02.aflexpod.cisco.com

3 nodes configured
4 resources configured

Online: [ SQLVM01.aflexpod.cisco.com SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]

Full list of resources:

Master/Slave Set: LINUX_SQLAG-master [LINUX_SQLAG]
Masters: [ SQLVM01.aflexpod.cisco.com ]
Slaves: [ SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]
SQLAG1_LISTNER_VIP (ocf::heartbeat:IPaddr2): Started SQLVM01.aflexpod.cisco.com

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
[root@SQLVM01 ~]# shutdown now
```

```
C:\Users\rgopunar>ping SQLAG1_LISTNER -t
Pinging SQLAG1_LISTNER.aflexpod.cisco.com [192.168.94.55] with 32 bytes of data:
Reply from 192.168.94.55: bytes=32 time<1ms TTL=63
Reply from 192.168.94.55: bytes=32 time<1ms TTL=63
Reply from 192.168.94.55: bytes=32 time<1ms TTL=63
Request timed out.
Reply from 192.168.94.55: bytes=32 time<1ms TTL=63
```

The following figure shows that the Availability Group resources moved to SQLVM02 replica.

Figure 185 Automatic Failover of AG Resources

```
[root@SQLVM02 ~]# pcs status
Cluster name: sqlavg-clus
Stack: corosync
Current DC: SQLVM02.aflexpod.cisco.com (version 1.1.16-12.e17-94ff4df) - partition with quorum
Last updated: Mon Jan 7 01:22:27 2019
Last change: Mon Jan 7 00:56:11 2019 by root via crm_resource on SQLVM02.aflexpod.cisco.com

3 nodes configured
4 resources configured

Online: [ SQLVM02.aflexpod.cisco.com SQLVM03.aflexpod.cisco.com ]
OFFLINE: [ SQLVM01.aflexpod.cisco.com ]

Full list of resources:

Master/Slave Set: LINUX_SQLAG-master [LINUX_SQLAG]
Masters: [ SQLVM02.aflexpod.cisco.com ]
Slaves: [ SQLVM03.aflexpod.cisco.com ]
Stopped: [ SQLVM01.aflexpod.cisco.com ]
SQLAG1_LISTNER_VIP (ocf::heartbeat:IPaddr2): Started SQLVM02.aflexpod.cisco.com

Daemon Status:
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
[root@SQLVM02 ~]#
```

# FlexPod Management Tools Setup

## Deploy NetApp Virtual Storage Console 7.2.1

This section describes the deployment process for the NetApp Virtual Storage Console (VSC).

### Virtual Storage Console 7.2.1 Pre-installation Considerations


The following licenses are required for VSC, on storage systems that run ONTAP 9.5:

- Protocol licenses (NFS and iSCSI)
- NetApp FlexClone (for provisioning and cloning only)
- NetApp SnapRestore (for backup and recovery)
- The NetApp SnapManager Suite


### Install Virtual Storage Console 7.2.1

To install the VSC 7.2.1 software by using an Open Virtualization Format (OVF) deployment, follow these steps:


1. Go to vSphere Web Client > Host Cluster > Deploy OVF Template.



2. Browse to the VSC OVF file downloaded from the NetApp Support site.




3. Enter the virtual machine name and select a datacenter or folder in which to deploy. Click Next.




4. Select a host cluster resource in which to deploy OVF. Click Next.

5. Review the details and accept the license agreement.


6. Select Storage.




7. From Select Networks, select a destination network and click Next.




8. From Customize Template, enter the VSC administrator password.




9. Select vCenter Registration Configuration.



10. Select Network Properties and customize the settings. Click Next.



11. Click Finish to complete the deployment of NetApp VSC virtual machine.



12. Power on the NetAppVSC virtual machine and open Console.
13. During the NetAppVSC virtual machine boot process, you see a prompt to install VMware Tools. From vCenter, select NetAppVSC VM > Guest OS > Install VMware Tools.
14. Networking configuration and vCenter registration information was provided during the OVF template customization; so after NetAppVSC VM is running, VSC, vSphere API for Storage Awareness (VASA), and VMware Storage Replication Adapter (SRA) are registered with vCenter.
15. Log out of the vCenter Client and log in again. From the Home menu, confirm that the NetApp VSC is installed.
16. If the VSC is not registered with vCenter, follow the steps in next section to manually register VSC with vCenter Server.

### Register Virtual Storage Console with vCenter Server

To register the VSC with vCenter Server, follow these steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open <https://localhost:8143/Register.html> in Internet Explorer.
2. Click Continue to This Website (Not Recommended).
3. In the Plug-in Service Information section, select the local IP address of the VSC virtual machine.
4. In the vCenter Server Information section, enter the host name or IP address, the user name (FlexPod admin user or root), and the user password for the vCenter Server. Click Register to complete the registration.

#### vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.


<b>Plugin service information</b>	
Host name or IP Address:	<input type="text" value="vsc.vikings.cisco.com"/>
<b>vCenter Server information</b>	
Host name or IP Address:	<input type="text" value="vc.vikings.cisco.com"/>
Port:	<input type="text" value="443"/>
User name:	<input type="text" value="administrator@vsphere.local"/>
User password:	<input type="password" value="••••••••"/>

**Register**


## Enable VASA and SRA

To enable VASA and SRA, follow these steps:


1. With a successful registration, the storage controller discovery automatically begins.
2. Open up the VSC Plug-in GUI.



3. Go to Configuration > Manage Extensions > Enable VASA and SRA.




4. Select the Enable options for VASA and SRA. Click Apply.



## Discover and Add Storage Resources

To discover the storage resources for the Monitoring and Host Configuration capability and the Provisioning and Cloning capability, follow these steps:

1. Using the vSphere web client, log into the vCenter Server as the FlexPod admin user. If the vSphere web client was previously opened, close it and reopen it.
2. On the Home screen, select the Home tab and click Virtual Storage Console.
3. Select Storage Systems. Some storage systems are shown as Unknown with their IP addresses.
4. From the Objects tab, click the IP address of an ONTAP Storage Cluster. Select Actions > Modify.



5. In the IP Address/Hostname field, enter the storage cluster management IP address. Enter admin for the user name and the admin password for password. Confirm Use TLS to Connect to This Storage System is selected. Click OK.

**Modify Storage System - Unknown - 192.168.166.20**

IP Address/Hostname: \* 192.168.166.20

User name: \* admin

Password: \*\*\*\*\*

Use TLS to connect to this storage system

Port: \* 443

Skip monitoring of this storage system

OK Cancel

6. Click OK to accept the controller privileges.

**Privileges**

**Allowed Privileges**

Create Storage	This role allows for the creation of volumes and logical unit numbers (LUNs).
Discovery	This role allows for the discovery of all the connected storage controllers.
Destroy Storage	This role allows for the destruction of volumes and LUNs. Includes all the privileges from Create Storage and Modify Storage.
Modify Storage	This role allows for the resizing and deduplicating of storage. Includes all the privileges from Create Storage.
PBM	This role allows for policy-based management of storage using storage capabilities.

**Disallowed Privileges**

OK Cancel


7. Wait for the Storage Systems to update. You might need to click Refresh to complete this update. The Storage Systems names display for the Storage Cluster and Storage Virtual Machines on the cluster for which the admin credentials are entered.

8. Enter the user credentials for all storage clusters and refresh.

## Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, follow these steps:

1. From the Home screen, click vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for these hosts.




2. Check the settings to apply to the selected vSphere hosts. Click OK to apply the settings.

---

 This functionality sets values for HBAs and converged network adapters (CNAs), sets appropriate paths and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

---

3. Click OK.



4. From the Home screen in the vSphere Web Client, select Virtual Storage Console.
5. On the left under Virtual Storage Console, select NFS VAAI Tools.
6. Make sure that NFS Plug-in for VMware VAAI Version 1.1.2-3 is shown.
7. Click Install on Host.
8. Select both ESXi hosts and click Install.

9. For each host for which settings were adjusted in the previous step, place the host in maintenance mode, reboot the host, and exit maintenance mode.



In testing, a conflict was identified between NetApp VSC, vSphere 6.7U1, and UCS servers containing the LSI MegaRAID SAS Invader local disk controller. If NetApp VSC cannot discover and set optimized settings for a server, the server most likely has the MegaRAID SAS Invader controller installed. This controller uses the lsi-mr3 ESXi vib. A workaround for this problem is to disable the disk controller by running `esxcli system module set --enabled=false --module=lsi_mr3` from an ESXi console or SSH prompt and then rebooting the host. When the host comes back up, the NetApp VSC functions should then work.

---


## Virtual Storage Console 7.2.1 Provisioning Datastores

Using VSC, you can provision the NFS, FC or iSCSI datastore and attach it to a single host or multiple hosts in the cluster. The following sections explain how to provision a datastore and attach it to the cluster.


### Provision NFS Datastore

To provision the NFS datastore, follow these steps:


1. From the Home screen of the vSphere Web Client, right-click the FlexPod-Management cluster and select NetApp VSC > Provision Datastore.




2. Enter the datastore name and select the type NFS.
3. Click Next.




- 4. Select the cluster name in the storage system and desired SVM to create the datastore. In this example, Infra-SVM is selected.
- 5. Click Next.



6. Enter the size of the datastore and select the aggregate name.
7. Click Next.



8. Review the details and click Finish.



9. Click OK.



The datastore is created and mounted on all hosts in the cluster. Click Refresh from the vSphere web client to see the newly created datastore.

## Provision iSCSI Datastore

To provision the iSCSI datastore, follow these steps:

1. From the Home screen of the vSphere Web Client, right-click the FlexPod Management cluster and select NetApp VSC > Provision Datastore.
2. Enter the datastore name and select the type as VMFS. For VMFS protocol, select iSCSI.
3. Click Next.

**NetApp Datastore Provisioning Wizard**

1 Name and type  
2 Storage system  
3 Details  
4 Ready to complete

Specify the name and type of datastore you want to provision.

You will be able to select the storage system for your datastore in the next page of this wizard.

Name : \* infra\_iSCSI\_datastore\_1

Type : \*  NFS  VMFS

VMFS Protocol : \*  FC/FCoE  iSCSI

Back Next Finish Cancel

4. Select the cluster name in the storage system and the desired SVM to create the datastore. In this example, Infra-SVM is selected.
5. Click Next.


The screenshot shows the 'NetApp Datastore Provisioning Wizard' window. The left sidebar indicates the current step is '2 Storage system'. The main area is titled 'Select the storage system you want to use to provision new datastore.' Below this, it states 'The list of storage systems below is filtered by the datastore type and protocol information entered on the previous page.' The 'vCenter Server' is set to 'vc.vikings.cisco.com'. The 'Storage system' dropdown is set to 'bb04-affa300 (192.168.156.60)'. The 'SVM' dropdown is set to 'Infra-SVM'. There is a section for 'Unusable storage systems' which is currently empty. At the bottom, there are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

6. Enter the size of the datastore. Select the Create new volume check box and select the aggregate name.

7. Click Next.

The screenshot shows the 'NetApp Datastore Provisioning Wizard' window at step '3 Details'. The left sidebar shows '3 Details' as the active step. The main area is titled 'Specify details of the new datastore.' It includes a checkbox for 'Thin provision' which is unchecked. The 'Size (GB)' field is set to '100'. The 'Create new volume' checkbox is checked. The 'Aggregate' dropdown is set to 'aggr1\_node02 - ( 6377.6 GB - Free )'. The 'Datastore cluster' dropdown is set to 'None'. At the bottom, there are 'Back', 'Next', 'Finish', and 'Cancel' buttons.

8. Review the details and click Finish.



9. Click OK.



The datastore is created and mounted on all the hosts in the cluster. Click the Refresh screen from vSphere web client to see the newly created datastore.

## Deploy NetApp SnapCenter

The SnapCenter platform is based on a multitier architecture that includes a centralized management server (SnapCenter Server) and a SnapCenter host agent.

### SnapCenter Server Requirements

Table 14 lists the minimum requirements for installing SnapCenter Server and plug-in on a Windows server. For the latest version compatibility and other plug-in information, go to the [NetApp Interoperability Matrix Tool](#).

Refer [SnapCenter 4.1.1 Installation and Setup Guide](#) for complete documentation.

**Table 14 SnapCenter Server Requirements**


Component	Requirements
Minimum CPU count	4 cores/vCPUs
Memory	Minimum: 8GB Recommended: 32GB
Storage space	Minimum space for installation: 10GB Minimum space for repository: 20GB
Supported operating systems	Windows Server 2012

	Windows Server 2012 R2 Windows Server 2016
Software packages	.NET 4.5.2 or later Windows Management Framework 4.0 or later PowerShell 4.0 or later Java 1.8 (64-bit)
Active Directory domain membership	Windows Host must be joined to AD domain.
Database for SnapCenter repository	MySQL Server 5.7.22 (installed as part of the SnapCenter installation)

**Table 15 Port Requirements**

Port	Requirement
443	vCenter Server to SnapCenter Server API Access over HTTPS
8144	SnapCenter GUI to SnapCenter Plug-in for VMWare
8145	SnapCenter Server Core API
8146	For SnapCenter server REST API
3306	MySQL
443	vCenter Server to SnapCenter Server API Access over HTTPS

Figure 186 NetApp Snapshot Deployment



### SnapCenter 4.1.1 License Requirements

The following licenses are required for SnapCenter, on storage systems that run ONTAP 9.5. These licenses do not need to be added into the SnapCenter GUI:

- Protocol licenses (NFS, FCP or iSCSI as per protocol selection for deployment)
- NetApp FlexClone (for provisioning and cloning only)
- NetApp SnapRestore (for backup and recovery)
- The NetApp SnapManager Suite

### SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, APIs, and the SnapCenter repository. SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You can accomplish high availability by using (NLB) and Application Request Routing (ARR) with SnapCenter. For larger environments with thousands of hosts, adding multiple SnapCenterServers can help balance the load.

### SnapCenter Plug-in for VMware vSphere

The Plug-in for VMware vSphere is a host-side component of the NetApp storage solution. It provides a vSphere web client GUI on vCenter to protect VMware virtual machines and datastores, and supports SnapCenter application-specific plug-ins in protecting virtualized databases and file systems.

## Support for Virtualized Databases and File Systems

The Plug-in for VMware vSphere provides native backup, recovery, and cloning of virtualized applications (virtualized SQL and Oracle databases and Windows file systems) when using the SnapCenter GUI.

SnapCenter natively leverages the Plug-in for VMware vSphere for all SQL, Oracle, and Windows file system data protection operations on virtual machine disks (VMDKs), raw device mappings (RDMs), and NFS datastores.

## SnapCenter Installation

Before you install SnapCenter, the Your SnapCenter Server host system must be up-to-date with Windows updates with no pending system restarts.

1. Download the SnapCenter Server installation package from [NetApp Support](#) site.
2. Install the SnapCenter Server by double-clicking the downloadable .exe file to launch the SnapCenter Server installer.
3. On the Prerequisites Validation screen, the host is validated to see if it meets the requirements to install the SnapCenter Server. If the minimum requirements are not met appropriate error or warning messages are displayed. If the restart of the host is pending, a warning message is displayed.
4. In the Network Load Balancing screen, if you want to enable and configure NLB select Enable and configure NLB on the host.
  - a. Select Create New NLB Cluster and enter the details for first node.
  - b. After creating the NLB cluster on the first node, when you run the installer on a second node, select Join Existing NLB Cluster and enter the details.
5. On the Credentials screen, enter the credentials that you want to use to log in to SnapCenter as the administrator.
  - a. The SnapCenter Server web component and SnapCenter SMCORE Service are installed in the corresponding folders at the default location C:\Program Files\NetApp.
  - b. The repository component is installed at the default location C:\ProgramData\NetApp\SnapCenter.
6. On the SnapCenter Ports Configuration screen, enter the port details. The default ports are auto populated but you can specify a custom port. In a NLB setup for the second node, the ports used while installing on the first node are auto populated and are not configurable.
7. On the MySQL Database Connection screen, enter the MySQL database password.
8. On the Ready to Install screen, click Install. Log files are listed (oldest first) in the %temp% folder.

## SnapCenter Configuration

The default GUI URL is a secure connection to the default port 8146 on the server where the SnapCenter Server is installed (**https://server:8146**). If you provided a different server port during the SnapCenter installation, that port is used instead.

For NLB deployment, you must access SnapCenter using the NLB cluster IP (**https://NLB\_Cluster\_IP:8146**). If you do not see the SnapCenter UI when you navigate to **https://NLB\_Cluster\_IP:8146** in IE, you must add the NLB IP address as a trusted site in IE on each plug-in host, or you must disable IE Enhanced Security on each plug-in host.

## Add a User or Group to a Role

To add a user or group to a role, follow these steps:

1. In the left navigation pane, click Settings.
2. In the Settings page, click Roles.
3. In the Roles page, select the role to which you want to add the user.
4. Click Modify.
5. Click Next until you reach the Users/Groups page of the wizard.
6. Select Domain or Workgroup as the Authentication type. For Domain authentication type, you should specify the domain name of the user or group to which you want to add the user to a role.
7. Select either User or Group.
8. In the User or Group Name field, enter a user or group name, and then click Add.
9. Repeat this step to add additional users or groups to the selected role.
10. Click Next to view the summary, and then click Finish.

## Configure Storage System Connections

To perform data protection and provisioning operations with SnapCenter, you must first set up the storage system connections that give SnapCenter access to ONTAP storage. Storage Systems can also be added thru SnapCenter Plugin for vCenter. Both methods require the same set of information.

## Install SnapCenter Plug-in for VMware

### Host and Privilege Requirements for the Plug-in for VMware vSphere

Review the following requirements before you install the SnapCenter Plug-in for VMware vSphere


- You must have SnapCenter admin privileges to install and manage the SnapCenter GUI.
- You must install the Plug-in for VMware vSphere on a Windows host (virtual host or physical host). The Plug-in for VMware vSphere must be installed on a Windows host regardless of whether you use the plug-in to protect data on Windows systems or Linux systems.
- When installing a plug-in on a Windows host, if you specify a Run As account that is not built-in or if the Run As user belongs to a local workgroup user, you must disable UAC on the host.
- Do not install the Plug-in for VMware vSphere on the vCenter Server appliance.
- You must not install the Plug-in for VMware vSphere on the vCenter Server appliance, which is a Linux host. You can only install the Plug-in for VMware vSphere on Windows hosts.
- You must not install other plug-ins on the host on which the Plug-in for VMware vSphere is installed.
- You must install and register a separate, unique instance of the Plug-in for VMware vSphere for each vCenter Server.

- Each vCenter Server, whether or not it is in Linked mode, must be paired with a separate instance of the Plug-in for VMware vSphere.
- Each instance of the Plug-in for VMware vSphere must be installed on a separate Windows host. One instance can be installed on the SnapCenter Server host.
- vCenters in Linked mode must all be paired with the same SnapCenter Server.

For example, if you want to perform backups from six different instances of the vCenter Server, then you must install the Plug-in for VMware vSphere on six hosts (one host can be the SnapCenter Server host) and each vCenter Server must be paired with a unique instance of the Plugin for VMware vSphere.

## Run As Credentials

Before you can perform data protection operations, you must set up the storage virtual machine connections and add Run As credentials that the SnapCenter Server and the SnapCenter plug-ins use.



1. Domain administrator or any member of the administrator group. Specify the domain administrator or any member of the administrator group on the system on which you are installing the SnapCenter plug-in. Valid formats for the **Username** field are:

```
NetBIOS\UserName
Domain FQDN\UserName
UserName@upn
```

2. Local administrator (for workgroups only). For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the **Username field is: UserName**
3. Run As credentials for individual resource groups. If you set up Run As credentials for individual resource groups and the username does not have full admin privileges, you must assign at least the resource group and backup privileges to the username.

To perform data protection and provisioning operations with SnapCenter, first you must set up the storage system connections using SnapCenter Plug-in for vCenter GUI that, to give SnapCenter access to ONTAP storage as described in next section.

## Install the Plug-in for VMware vSphere from the SnapCenter GUI


To install the plug-in for VMware vSphere from the SnapCenter GUI, follow these steps:

1. In the left navigation pane, click Hosts. Verify that Managed Hosts is selected at the top.
2. Click Add.
3. On the Hosts page, specify:
  - a. Host OS: vSphere
  - b. Host Name: Host on which you want to install the Plug-in for VMware
  - c. Run As Name: Run As account configured with user credentials with access to the host
  - d. Port: Leave default 8145
4. On the Plug-ins to install page, specify the vCenter information
5. Review the summary, and then click Finish
6. SnapCenter Server performs the following tasks during the installation process:
  - a. Adds the host to the SnapCenter registry.
  - b. Installs the Plug-in for VMware vSphere, if the plug-in is not already installed on the host.
  - c. If the Plug-in for VMware vSphere is installed on the SnapCenter Server host, it also installs the SnapCenter Plug-in for Microsoft Windows to support SnapCenter repository backup operations by using PowerShell cmdlets. You cannot use any other feature of this instance of the Plug-in for Windows.
  - d. Adds the Plug-in for VMware vSphere web client to vCenter.

## Configure SnapCenter Plug-in for vCenter

After you have successfully installed the Plug-in for VMware vSphere, to configure SnapCenter to make it ready to backup virtual machines, follow these steps:


1. In your browser, navigate to VMware vSphere Web Client URL <https://<vCenter Server>/vsphere-client/?csp>.
2. On the VMware vCenter Single Sign-On page, log on.
3. On the VMware vSphere Web Client page, click Home and select SnapCenter Plug-in for VMware vSphere to Bring up SnapCenter Plug-in for VMware GUI. It will take you to the SnapCenter Plug-in for VMware Dashboard.




### Add Storage Systems (Storage Virtual Machine)

To add storage systems, follow these steps:


1. Go to the Storage Systems page.



2. Click + Add Storage System to add an SVM.




3. Enter vCenter, Storage System, user credentials and other required information in following Dialog.



## Create Backup Policies for Virtual Machines and Datastores

To create backup policies for virtual machines and datastores, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Policies.



2. On the Policies page, click + New Policy in the toolbar.
3. On the New Backup Policy page:
  - a. Enter the policy name and a description.
  - b. Enter backups to keep.
  - c. From the Frequency drop-down list, select the backup frequency (hourly, daily, weekly, monthly, and on-demand only).
  - d. Click Add.

**New Backup Policy**

**Name**

**Description**

**vCenter Server**

**Retention** Days to keep

**Frequency**

**Replication**

Update SnapMirror after backup

Update SnapVault after backup

Snapshot label

**Advanced**

VM consistency

Include datastores with independent disks

**Scripts**

Cancel Add

4. Create multiple policies as required for different sets of virtual machines or datastores.

### Create Resource Groups


Resource groups are groups of virtual machines or datastores that are backed up together. A backup policy is associated with the resource group to back up the resource are required frequency and retain certain number of backups as defined in the policy.

To create resource groups, follow these steps:

1. In the left Navigator pane of the VMware vSphere web client, click Resource Groups and then click Create Resource Group. This is the easiest way to create a resource group. However, you can also create a resource group with one resource by performing one of the following steps:
  - a. To create a resource group for one virtual machine, click VMs and Templates then right-click a virtual machine, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.
  - b. To create a resource group for one datastore, click Storage then right-click a datastore, select NetApp SnapCenter from the drop-down list, and then select Create Resource Group from the secondary drop-down list.




2. On the General Info & Notification page, enter the resource group name and complete the notification settings. Click Next.




3. Select a datastore as parent entity to create a resource group of virtual machines, and then select the virtual machines from the available list. Click Next.





7. Review the Summary and click Finish to complete the creation of the resource group.




## View Virtual Machine Backups and Restore from vCenter by Using SnapCenter Plug-in


### View Backups

Backups of the resources included in the resource group occurs per the schedule of the policies associated with the resource group. To view the backups created per schedule, complete one of the following methods:


1. Go to any virtual machine that has been added to Resource Group > More Objects > Backups. It shows all the backups for that resource.



2. On the VMware vSphere Web Client page, click Home and select SnapCenter Plug-in for VMware vSphere to Bring up SnapCenter Plug-in for VMware GUI. Click Resource Groups and select any resource group. In the right pane, the completed backups are displayed.




3. An on-demand backup can also be created for any resource group by following these steps:
  - a. Click the virtual machine or datastore resource > More Objects > Resource Groups. Select the Resource Group and click Run Now. This process triggers the backup of the resource group containing a particular resource.




## Restore from vCenter by Using SnapCenter Plug-in

To restore from vCenter by using SnapCenter Plug-in, follow these steps:

1. The Plug-in for VMware vSphere provides native backup, recovery, and cloning of virtualized applications. Go to vCenter Web Client GUI > Select a virtual machine > right-click and select NetApp SnapCenter.




2. Select a backup from which to restore. Click Next.




3. From the Restore Scope drop-down list, select either Entire virtual machine to restore the virtual machine with all VMDKs or select Particular Virtual Disk to restore the VMDK without affecting the virtual machine configuration and other VMDKs.





5. Review the Summary and click Finish to complete the restore process.



## Summary

---

FlexPod is the optimal shared infrastructure foundation to deploy a variety of IT workloads. It is built on leading computing, networking, storage, and infrastructure software components. This CVD is a detailed guide for deploying Microsoft SQL 2017 deployment on Linux virtual machine deployed either in VMware and Hyper-V virtual environments. This document also explains the process for deploying SQL Always On Availability Groups feature (using pacemaker cluster resource manager), which helps to achieve application-level high availability in addition to FlexPod's robust infrastructure-level high availability.

## Appendix

---

### FlexPod Backups

#### Cisco UCS Backup

Automated backup of the Cisco UCS domain is important for recovery of the Cisco UCS Domain from issues ranging from catastrophic failure to human error. There is a native backup solution within Cisco UCS that allows local or remote backup using FTP/TFTP/SCP/SFTP as options and is detailed below.

Backups created can be a binary file containing the Full State, which can be used for a restore to the original or a replacement pair of fabric interconnects. Alternately this XML configuration file consisting of All configurations, just System configurations, or just Logical configurations of the Cisco UCS Domain. For scheduled backups, options will be Full State or All Configuration, backup of just the System or Logical configurations can be manually initiated.

To schedule the backup, follow these steps within the Cisco UCS Manager GUI:

1. Select Admin within the Navigation pane and select All.
2. Click the Policy Backup & Export tab within All.
3. For a Full State Backup, All Configuration Backup, or both, specify the following:
  - a. Hostname : <IP or FQDN of host that will receive the backup>
  - b. Protocol: [FTP/TFTP/SCP/SFTP]
  - c. User: <account on host to authenticate>
  - d. Password: <password for account on host>
  - e. Remote File: <full path and filename prefix for backup file>
  - f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on save>
  - g. Schedule: [Daily/Weekly/Bi Weekly]

Figure 187 UCS Configuration Backup

All

General Policy Backup & Export

---

Full State Backup Policy

Hostname : 192.168.166.150

Protocol :  FTP  TFTP  SCP  SFTP

User : root

Password : \*\*\*\*\*

Remote File : /var/www/html/configs/ucs/6332.full

Admin State :  Disable  Enable

Schedule :  Daily  Weekly  Bi Weekly

Max Files : 0

Description : Database Backup Policy

---

All Configuration Backup Policy

Hostname : 192.168.166.150

Protocol :  FTP  TFTP  SCP  SFTP

User : root

Password : \*\*\*\*\*

Remote File : /var/www/html/configs/ucs/6332.config

Admin State :  Disable  Enable

Schedule :  Daily  Weekly  Bi Weekly

Max Files : 0

Description : Configuration Export Policy

---

Backup/Export Config Reminder

Admin State :  Disable  Enable

## Cisco Nexus Backups

The configuration of the Cisco Nexus 9000 switches can be backed up manually at any time with the copy command, but automated backups can be put in place with the NX-OS feature scheduler. An example of setting up automated configuration backups of one of the Cisco Nexus 93180YC-EX switches is shown below:

```
bb09-93180-a# conf t
```

Enter configuration commands, one per line. End with CNTL/Z.

```
bb09-93180-a(config)# feature scheduler
```

```
bb09-93180-a(config)# scheduler logfile size 1024
```

```
bb09-93180-a(config)# scheduler job name backup-cfg
```

```
bb09-93180-a(config-job)# copy running-config
tftp://192.168.166.150/93180/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
```

```
bb09-93180-a(config-job)# exit
```

```
bb09-93180-a(config)# scheduler schedule name daily
```

```
bb09-93180-a(config-schedule)# job name backup-cfg
```

```
bb09-93180-a(config-schedule)# time daily 2:00
```

```
bb09-93180-a (config-schedule) # end
```

For detailed information, go to: [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system\\_management/configuration/guide/b\\_Cisco\\_Nexus\\_9000\\_Series\\_NX-OS\\_System\\_Management\\_Configuration\\_Guide\\_7x/b\\_Cisco\\_Nexus\\_9000\\_Series\\_NX-OS\\_System\\_Management\\_Configuration\\_Guide\\_7x\\_chapter\\_01010.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01010.html)

## About the Authors

---

Gopu Narasimha Reddy, Technical Marketing Engineer, Compute Systems Product Group, Cisco Systems, Inc.

Gopu Narasimha Reddy is a Technical Marketing Engineer working with Cisco UCS Datacenter Solutions group. His current focus is to develop, test and validate solutions on Cisco UCS platform for Microsoft SQL Server databases on Microsoft Windows and VMware platforms. He is also involved in publishing TPC-H database benchmarks on Cisco UCS servers. His areas of interest include building and validating reference architectures, development of sizing tools in addition to assisting customers in SQL deployments.

Sanjeev Naldurgkar, Technical Marketing Engineer, Compute Systems Product Group, Cisco Systems, Inc.

Sanjeev has been with Cisco for six years focusing on delivering customer-driven solutions on Microsoft Hyper-V and VMware vSphere. He has over 16 years of experience in the IT Infrastructure, Server virtualization, and Cloud Computing. He holds a Bachelor's Degree in Electronics and Communications Engineering, and leading industry certifications from Microsoft and VMware.

Atul Bhalodia, Sr. Technical Marketing Engineer, NetApp Cloud Infrastructure Engineering, NetApp

Atul Bhalodia is a Sr. Technical Marketing Engineer in the NetApp and Cloud Infrastructure Engineering team. He focuses on the Architecture, Deployment, Validation, and documentation of cloud infrastructure solutions that include NetApp products. Prior to his current role, he was a Software Development Lead for NetApp SnapDrive and SnapManager Products. Atul has worked in the IT industry for more than 20 years and he holds Master's degree in Computer Engineering from California State University, San Jose, CA.

## Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- John George, Cisco Systems, Inc.